arXiv:1605.08122v1 [math.PR] 26 May 2016

ON THE MIXING TIME OF KAC'S WALK AND OTHER HIGH-DIMENSIONAL GIBBS SAMPLERS WITH CONSTRAINTS

NATESH S. PILLAI[‡] AND AARON SMITH[♯]

Abstract. Determining the total variation mixing time of Kac's random walk on the special orthogonal group $\mathrm{SO}(n)$ has been a long-standing open problem. In this paper, we construct a novel non-Markovian coupling for bounding this mixing time. The analysis of our coupling entails controlling the smallest singular value of a certain random matrix with highly dependent entries. The dependence of the entries in our matrix makes it not-amenable to existing techniques in random matrix theory. To circumvent this difficulty, we extend some recent bounds on the smallest singular values of matrices with independent entries to our setting. These bounds imply that the mixing time of Kac's walk on the group $\mathrm{SO}(n)$ is between $C_1 n^2$ and $C_2 n^4 \log(n)$ for some explicit constants $0 < C_1, C_2 < \infty$, substantially improving on the bound of $O(n^5 \log(n)^2)$ in the preprint [Jia12b] of Jiang. Our methods may also be applied to other high dimensional Gibbs samplers with constraints and thus are of independent interest. In addition to giving analytical bounds on the mixing time, our approach allows us to compute rigorous estimates of the mixing time by simulating the eigenvalues of a random matrix.

1. Introduction

Mark Kac introduced a random walk on the sphere in his 1954 paper [Kac54] as a model for a Boltzmann gas. In this paper, we study Kac's walk on the special orthogonal group $\mathrm{SO}(n)$, which was first introduced in a statistical context [Has70] and has been studied as a generalization of Kac's walk on the sphere since [DSC00] (see also, *e.g.*, [Oli09, PS07, CCL03, Jan03, HJ16]).

Kac's walk on $\mathrm{SO}(n)$ is a discrete-time Markov chain $\{X_t\}_{t\geq 0}$ that evolves as follows. Fix an ordering of the $N \equiv \frac{n(n-1)}{2}$ planes generated by two coordinates in $\mathbb{R}^n$ and choose $X_0 \in \mathrm{SO}(n)$. For $t \in \mathbb{N}$, choose $1 \leq i_t \leq N$ and $\theta \in [0, 2\pi]$ uniformly at random and set

$$X_{t+1} = \mathrm{R}(i_t, \theta_t) X_t, \tag{1.1}$$

where $\mathrm{R}(i, \theta)$ denotes a rotation by the angle θ in the i'th coordinate plane. If the i'th coordinate plane is associated with the coordinate axes $1 \leq k < \ell \leq n$, $\mathrm{R}(i, \theta)$ is an $n \times n$ matrix with entries

$$\begin{aligned} \mathrm{R}(i,\theta)_{jj} &= \cos(\theta), && j \in \{k, \ell\} \\ \mathrm{R}(i,\theta)_{k\ell} &= \sin(\theta), & \mathrm{R}(i,\theta)_{\ell k} &= -\sin(\theta) \\ \mathrm{R}(i,\theta)_{jj} &= 1, && j \notin \{k, \ell\} \\ \mathrm{R}(i,\theta)_{jj'} &= 0, && j' \notin \{j, k, \ell\}. \end{aligned} \tag{1.2}$$

[‡]pillai@fas.harvard.edu, Department of Statistics, Harvard University, 1 Oxford Street, Cambridge MA 02138, USA.

[♯]smith.aaron.matthew@gmail.com, Department of Mathematics and Statistics, University of Ottawa, 585 King Edward Avenue, Ottawa ON K1N 7N5, Canada.

If we write $X_t = [v_t^{(1)}\, v_t^{(2)}\, \dots\, v_t^{(n)}]$, the law of $\{v_t^{(1)}\}_{t\geq 0}$ is known as *Kac's walk on the sphere* S^{n-1}. Physically, Kac motivated this random walk by considering n particles in a one-dimensional box. He assumed that these particles were uniformly distributed in space, and the vector $v_t^{(1)}$ models the change in their velocities over time as collisions occur; the condition that $v_t^{(1)}$ be constrained to the sphere corresponds to the principle of conservation of energy. Understanding the mixing properties of this process is central to Kac's program in kinetic theory (see [MM13] for a useful description of this program). Kac's walks on the sphere and on $\mathrm{SO}(n)$ have attracted great attention and estimating their mixing times has been a long standing open problem (see Sections 1.1 and 1.2). Recently, in [PS15], the authors of this paper obtained a matching upper bound and lower bound for the mixing time of Kac's walk in S^{n-1}, thus settling this problem upto a constant factor.

To state our main result, we recall some standard definitions. For measures ν_1, ν_2 on a measure space $(\Omega, \mathcal{F})$, the *total variation distance* between ν_1, ν_2 is given by

$$\|\nu_1 - \nu_2\|_{\mathrm{TV}} = \sup_{A\in\mathcal{F}} (\nu_1(A) - \nu_2(A)).$$

We denote the distribution of a random variable X by $\mathcal{L}(X)$ and write $X \sim \nu$ as a shorthand for $\mathcal{L}(X) = \nu$. For a Markov chain $\{X_t\}_{t\geq 0}$ with unique associated stationary distribution ν on state space Ω, we define the associated *mixing profile* by

$$\tau(\epsilon) = \inf\{t : \sup_{X_0 = x\in\Omega} \|\mathcal{L}(X_t) - \nu\|_{\mathrm{TV}} < \epsilon\}$$

and the *mixing time* by $\tau_{\mathrm{mix}} = \tau(0.25)$.

Let μ denote the normalized Haar measure on $\mathrm{SO}(n)$. Our main result is the following bound on the mixing time of Kac's walk on $\mathrm{SO}(n)$:

Theorem 1.1. *Let $\{X_t\}_{t\geq 0}$ be a copy of Kac's walk on $\mathrm{SO}(n)$. Then for all sequences $T = T(n) > 10^7\, n^4 \log(n)$,*

$$\limsup_{n\to\infty} \sup_{X_0 = x\in \mathrm{SO}(n)} \|\mathcal{L}(X_T) - \mu\|_{\mathrm{TV}} = 0, \tag{1.3}$$

and for all sequences $T = T(n) < N$,

$$\liminf_{n\to\infty} \sup_{X_0 = x\in \mathrm{SO}(n)} \|\mathcal{L}(X_T) - \mu\|_{\mathrm{TV}} = 1. \tag{1.4}$$

We have not tried to optimize the constant 10^7 appearing in Theorem 1.1.

1.1. **Motivations Outside of Physics.** Kac's random walk has been studied in a wide range of fields including computer science, statistics and numerical analysis. To our knowledge, the Markov chain that we call Kac's walk on $\mathrm{SO}(n)$ was initially proposed in [Has70] as a Gibbs sampler targetting the Haar measure on $\mathrm{SO}(n)$. The problem of sampling from Haar measure on $\mathrm{SO}(n)$ was motivated by [Jam55], but the walk itself has been suggested as a computationally efficient method for finding projections onto random small-dimensional subspaces [AC06]. Bounds on the mixing time of Kac's walk are required to check that this approach is, in fact, computationally efficient.

Our analysis of Kac's walk is also interesting as a worked example that belongs to several active areas of research. The Markov chains we study are a sequence of high-dimensional Gibbs samplers (see [CG92]). Despite three decades of extensive work in this area, there are few effective bounds on the mixing times of Gibbs samplers in high dimensions (see

[JH01, DKSC08] for an introduction to the large literature on this problem). Of the existing effective bounds on the mixing times of high dimensional Gibbs samplers on continuous state spaces, almost all target distributions with support equal to a union of quadrants of $\mathbb{R}^n$ (*e.g.,* [JH01]) or involve explicitly computing spectral information for the transition kernel (*e.g.,* the analyses [Ros94, Por96, HJ16] of other random walks on $\mathrm{SO}(n)$). Our analysis gives one of relatively few results for Gibbs samplers on a complicated sample space for which spectral information cannot easily be used. Some closely related papers are [Lov99, LV03] on Gibbs samplers on convex sets, as well as the papers [Smi14, Jia12a] directly motivated by the study of Kac's walk on $\mathrm{SO}(n)$. There is also a large number of papers studying the mixing time of Markov chains on groups [Dia88, Ros94, Fur02, SC04, HJ16] using Fourier analysis and representation theory. Unlike these papers, we do not use Fourier analytic tools and thus our methods are generalizable to other Gibbs samplers in $\mathbb{R}^n$.

1.2. **Previous Work.** The central question in the study of Kac's walk is to determine the speed at which it converges to equilibrium. This question is somewhat vague, as it does not specify the metric under which convergence is to be measured. Early work focused on proving that the spectral gap of the chain was large. In [DSC00], the authors showed that the spectral gap of the walk on $\mathrm{SO}(n)$ was at least order of n^{-3}. Janvresse first showed in [Jan01] that the spectral gap of the walk on the sphere was exactly on the order of n^{-1}. Janvresse also showed in [Jan03] that the spectral gap of the walk on $\mathrm{SO}(n)$ was on the order of n^{-2}. The exact spectral gap for both walks was found in [CCL03], and the full spectrum was computed in [Mas03]. Some of this work was generalized in [Cap08]. Although these bounds imply a convergence *rate* for Kac's walk in L^2, and a bound on the *distance* to stationarity in L^2 implies a bound on the total variance distance to stationarity, these bounds *do not* imply any bound at all on the total variation mixing time of Kac's walk. This is because, when $\mathcal{L}(X_0)$ is concentrated at a point, the initial L^2 distance to stationarity is not finite.

Later work has focused on stronger metrics for convergence or more demanding versions of the problem. In [CCR$^+$08], a very strong convergence condition as measured by entropy was discovered. These bounds, like the bounds relating to spectral gap, only imply convergence for sufficiently smooth initial distribution $\mathcal{L}(X_0)$. In this note, we focus on convergence bounds that do not depend on the initial distribution $\mathcal{L}(X_0)$. The first bound with this property was obtained [PS07], in which the authors showed a convergence time of order $O(n^{2.5})$ in the L^1 Wasserstein metric, and [Oli09] improved this bound to $O(n^2 \log(n))$ in the stronger L^2 Wasserstein metric. This latter bound is tight up to factors of $\log(n)$, and will be essential to our effort. Related Wasserstein bounds have also been found in [CF14] for several similar models. However, a mixing bound in the Wasserstein metric does not directly imply any mixing bound in the total variation metric.

Thus the bounds obtained in [PS07, Oli09], despite their strength, do not give any information at all about the mixing time in total variation distance. The first bound on convergence in total variation was on the order of 4^{n^2} steps, obtained by Diaconis and Saloff-Coste in [DSC00]. No progress was made on this problem until the recent unpublished work of Yunjiang Jiang [Jia12b], in which the author obtained a mixing bound of order $n^5 \log^2(n)$.

Theorem 1.1 of our paper also implies a mixing bound on the order of $n^4 \log(n)$ for Kac's walk on S^{n-1}. This improves upon all bounds prior to the present author's recent work [PS15], which shows matching upper and lower bounds on the order of $n \log(n)$ for this walk.

The papers [Ros94, Por96, HJ16, MM13, HM14] all study variants or projections of Kac's walk on $\mathrm{SO}(n)$.

1.3. **Our Contributions.** We have three main contributions: an order of magnitude improvement on the previous best bound for the convergence rate of Kac's walk in the strong total variation metric, new bounds on the smallest singular values of certain random matrices with dependent entries, and a general approach to bounding the mixing times of Gibbs samplers on spaces that are not 'rectangular.' Our method also gives a way to compute effective mixing bounds via simulation. We now give a broad overview of our approach, and its relationship to some previous work.

The upper bound of $O(n^4 \log(n))$ on the mixing time of Kac's walk is proved by using the popular *coupling* technique: we run two copies $\{X_t\}_{t\geq 0}$, $\{Y_t\}_{t\geq 0}$ of Kac's walk, and study the first time $\inf\{t \geq 0 : X_t = Y_t\}$ that they collide. Like many non-Markovian couplings (see, *e.g.*, [HV03, Smi14, CV14, PS15]), the main idea is to construct a coupling in two passes: an initial, Markovian coupling of 'most' of the randomness in the chain, followed by a very general coupling of some 'leftover' randomness. Our initial coupling is exactly the one constructed in [Oli09]. Under this coupling, two copies of Kac's walk mix in Wasserstein distance after after $O(n^2 \log(n))$ steps. Our contribution is in the construction and analysis of the second stage coupling.

The usual approach to converting a Wasserstein mixing bound for a high-dimensional Gibbs sampler to a total variation mixing bound is via a greedy coupling: one attempts to match more and more coordinates as time progresses. Unfortunately, this approach works poorly for Kac's walk on $\mathrm{SO}(n)$. Indeed, as the authors discuss in [PS15], the greedy approach does not even work in the simpler case of Kac's walk on the sphere. Instead, we first run a scaffold $(\hat{X}_t, \hat{Y}_t)$ according to a coupling from [Oli09]. We then construct an N-dimensional perturbation of this scaffold by adding a small amount of additional randomness at each of N time steps. The key point here is that it turns out to be easier to analyze the coupling probability of our two chains by using the N bits of randomness all at once, rather than analyzing the coupling probability of the individual coordinate at each step as done in [PS15]. Our approach is somewhat reminiscent of the 'sprinkling' strategy used in random graph theory [AKS82], which also involves coupling 'most' random variables in an intuitive way and then carefully analyzing a small amount of 'leftover' randomness.

Our analysis does not depend critically on any special properties of $\mathrm{SO}(n)$, and our non-Markovian coupling construction in Section 2.1 makes sense for any Gibbs sampler. The analysis of our coupling is also applicable for other constrained high dimensional Gibbs samplers. This is in stark contrast to our previous work [PS15] on Kac's walk on the sphere.

Our simple results on random matrices in Section 6 are novel, giving bounds on the smallest singular values of random matrices with significant dependence between entries. These bounds are closely related to the results of [FG13, FV15] on the smallest singular values of random matrices with independent entries, and give bounds that are qualitatively similar to [FG13].

In addition to giving asymptotic results on mixing times, our method allows us to numerically estimate the mixing time of Kac's walk for fixed n by simulating the eigenvalues of a certain random matrix. This is useful for those interested in knowing the mixing time of a particular Gibbs sampler, and will generally give sharper results than our mathematical analysis. Note that estimating the mixing time of a Markov chain in this way is not trivial

- *a priori* it is not obvious how to obtain *any* rigorous bounds on the mixing time of Kac's walk by a finite computation. See, *e.g.*, [CV06] for further discussion on simulated estimates of mixing times.

1.4. **Notation.** For S either a finite set or a subset of $\mathbb{R}^m$ with finite nonzero Lebesgue measure, we write $\mathrm{unif}(S)$ for the uniform probability measure on S. For functions $f, g : \mathbb{N} \mapsto \mathbb{R}$, we write $f = O(g)$ if $\limsup_{n\to\infty} \frac{|f(n)|}{|g(n)|} < \infty$. We also write $f = \Omega(g)$ if $g = O(f)$ and we write $f = \Theta(g)$ if both $f = O(g)$ and $f = \Omega(g)$. Finally, we write $f = o(g)$ if $\limsup_{n\to\infty} \frac{|f(n)|}{|g(n)|} = 0$. Unless otherwise noted, the terms inside of such 'big-O' notation should always be taken with respect to the problem dimension n.

For $x, y \in \mathbb{R}^m$, we write $\langle x, y\rangle$ for the Euclidean inner product and $\|x-y\|$ for the associated norm. We also use '0' as a shorthand for the vector $(0, 0, \ldots, 0) \in \mathbb{R}^m$. Denote by $M(n)$ the collection of $n \times n$ matrices with real valued entries. For $h \in M(n)$, let $h^\dagger$ denote its transpose. For a linear map $T : \mathbb{R}^k \mapsto \mathbb{R}^\ell$, define the operator norm:

$$\|T\|_{\mathrm{Op}} = \sup_{\|v\|=1} \|Tv\|.$$

For elements $x_1, x_2, \ldots$ of a noncommutative group, we use the convention:

$$\prod_{s=1}^{t} x_s \equiv x_1 x_2 \ldots x_{t-1} x_t.$$

We write ∂S for the boundary of any set S. For any smooth manifold $\mathcal{M}$ and any $x \in \mathcal{M}$, we denote by $T_x\mathcal{M}$ the tangent space of $\mathcal{M}$ at x. We will often identify $T_x\mathbb{R}^m$ with $\mathbb{R}^m$. For any pair of smooth manifolds $\mathcal{M}$, $\mathcal{N}$ and any smooth function $f : \mathcal{M} \mapsto \mathcal{N}$, we define for $p \in \mathcal{M}$ the usual associated derivative map $\mathrm{d}f_p : T_p\mathcal{M} \to T_{f(p)}\mathcal{N}$. We recall an explicit construction of this map. Fix $v \in T_p\mathcal{M}$ and let $\gamma : [0, 1] \to \mathcal{M}$ satisfy $\gamma(0) = p$, $\gamma'(0) = v$. Then

$$\mathrm{d}f_p(v) = (f \circ \gamma)'(0).$$

The quantity $\mathrm{d}f_p(v)$ is independent of the path γ as long as $\gamma'(0) = v$. The rank of the linear map $\mathrm{d}f_p$ is denoted by $\mathrm{Rank}(\mathrm{d}f_p)$.

Let G be a Lie group G with Lie algebra $\mathcal{G}$. For $a \in G$, let $\mathrm{L}_a : G \mapsto G$ be the left multiplication map

$$\mathrm{L}_a(g) = ag.$$

The exponential map, denoted by exp, maps $\mathcal{G}$ onto G. When G is a matrix group and $\mathcal{G}$ is identified with a subset of $M(n)$, the exponential map has the explicit form

$$\exp(A) = \sum_{i=0}^{\infty} \frac{A^i}{i!}, \quad A \in \mathcal{G} \tag{1.5}$$

(see Section 1.1 of [AM]). The Hilbert-Schmidt inner product on $M(n)$ is:

$$\langle A, B\rangle_{\mathrm{HS}} = \mathrm{Tr}(A^\dagger B)$$

where Tr is the trace. The corresponding inner product is denoted by $\langle \cdot, \cdot \rangle_{\mathrm{HS}}$.

Throughout the paper, we will use the convention that the addition of angles θ, θ' is always done modulo 2π, and that the distance between two angles is measured with respect to the usual metric on the torus rather than the usual metric on the line.

The Lie algebra $\mathcal{G} = \mathfrak{so}(n)$ of $\mathrm{SO}(n)$ is the set of $n \times n$ skew-symmetric matrices

$$\mathfrak{so}(n) = \{h \in M(n) \,:\, h = -h^{\dagger}\}.$$

We denote by D_{HS} the Riemannian metric of $\mathrm{SO}(n)$ induced by the inner product $\langle \cdot, \cdot \rangle_{\mathrm{HS}}$ on $\mathfrak{so}(n)$. The Haar measure π on $\mathrm{SO}(n)$ is also induced by the inner product $\langle \cdot, \cdot \rangle_{\mathrm{HS}}$. We denote by $\mathcal{P} : \mathrm{SO}(n) \mapsto \mathfrak{so}(n)$ the orthogonal projection operator into $\mathfrak{so}(n)$ with respect to the Hilbert-Schmidt norm, *i.e.*,

$$\mathcal{P}(g) = \arg\min_{A \in \mathfrak{so}(n)} \|g - A\|_{\mathrm{HS}}.$$

We construct an orthonormal basis for $\mathfrak{so}(n)$ as follows. For $1 \leq i \leq N$, $\mathrm{R}_i(x) \equiv \mathrm{R}(i, x)$ (see Equation 1.1) is a map from the N-dimensional torus $\mathbb{T} = [0, 2\pi)^N$ to $\mathrm{SO}(n)$. Set

$$a_i = \frac{1}{\sqrt{2}} \mathrm{dR}_i(0) \in \mathfrak{so}(n).$$

The set $\{a_i\}_{1 \leq i \leq N}$ constitute an orthonormal basis in $\mathfrak{so}(n)$. This set of basis vectors were used in [Oli09] to obtain a coupling of two copies of random Kac's random walk on $\mathrm{SO}(n)$.

Throughout the paper, the quantities ϕ_n, ϵ_n and ω_n will control the three distance scales that are key to our coupling proof. The quantity ϕ_n, defined in Equation (5.1), will satisfy $\phi_n \leq (2n)^{-20}$ and controls the scale on which a certain function 'looks flat.' The quantity $\epsilon_n \leq \phi_n^{30}$ will control the total amount of 'injected randomness' available to our coupling. Finally, $\omega_n = \epsilon_n^{30}$ controls the typical distance between the two Markov chains that we are trying to couple.

2. Proof Strategy

We give an overview of the proof of Theorem 1.1. A basic ingredient in the proof of the upper bound on the mixing time in Theorem 1.1 is the following standard coupling inequality:

Lemma 2.1. *Let K be the transition kernel of a Markov chain with unique stationary distribution ν on state space Ω. Let $\{X_t\}_{t \geq 0}$, $\{Y_t\}_{t \geq 0}$ be two Markov chains, with $Y_0 \sim \nu$ started at stationarity and $X_0 = x \in \Omega$. Define the coalescence time*

$$\tau(x) = \inf\{t \,:\, X_t = Y_t\}. \tag{2.1}$$

Assume that the coupling of $\{X_t\}_{t \geq 0}$, $\{Y_t\}_{t \geq 0}$ satisfies $X_t = Y_t$ for all $t \geq \tau(x)$. Then

$$\|\mathcal{L}(X_t) - \nu\|_{\mathrm{TV}} \leq \mathbb{P}[\tau(x) > t].$$

Let $\{X_t\}_{t \geq 0}$, $\{Y_t\}_{t \geq 0}$ be two copies of Kac's walk as described in Equation (1.1). Let $\{i_t(x), \theta_t(x)\}_{t \geq 0}$ and $\{i_t(y), \theta_t(y)\}_{t \geq 0}$ be the random variables that determine the evolution of $\{X_t\}_{t \geq 0}$ and $\{Y_t\}_{t \geq 0}$ in Equation (1.1); we refer to such sequences as *update sequences*. To construct a coupling of $\{X_t\}_{t \geq 0}$ and $\{Y_t\}_{t \geq 0}$, it is enough to construct a coupling of $\{i_t(x), \theta_t(x)\}_{t \geq 0}$ and $\{i_t(y), \theta_t(y)\}_{t \geq 0}$.

The following coupling was defined and analyzed in [Oli09].

Definition 2.2 (Locally Contractive Coupling [Oli09])**.** For any fixed $T \in \mathbb{N}$, the following algorithm gives a coupling of $\{X_t, Y_t\}_{t=0}^{T}$ according to their associated update sequences:

(1) For $0 \leq t < T$, choose $i_t \sim \mathrm{unif}\{1, 2, \ldots, N\}$ and set

$$i_t(y) = i_t(x) = i_t.$$

(2) Let $\{\eta_t(x)\}_{0\le t<T}$ be an i.i.d. sequence of unif$[0,2\pi)$ random variables. We now construct the sequences $\{X_t, Y_t, h_t, \eta_t(y)\}_{t\ge 0}$ inductively in t. For $1 \le s \le t$, define

$$X_s = \Big(\prod_{u=s-1}^{0} \mathrm{R}(i_u(x), \eta_u(x))\Big) X_0, \;\; Y_s = \Big(\prod_{u=s-1}^{0} \mathrm{R}(i_u(y), \eta_u(y))\Big) Y_0.$$

Finally, set

$$h_s = \mathcal{P}(Y_s X_s^\dagger - \mathrm{Id}) \tag{2.2}$$
$$\eta_s(y) = \eta_s(x) + \frac{1}{\sqrt{2}} \langle h_s, a_{i_s} \rangle_{\mathrm{HS}}.$$

(3) Set $\theta_t(x) = \eta_t(x)$, $\theta_t(y) = \eta_t(y)$ for all $0 \le t < T$.

Theorem 1 of [Oli09] implies that it is possible to couple[1] two copies of Kac's random walk $\{X_t\}_{t\ge 0}$ and $\{Y_t\}_{t\ge 0}$ so that, for any $0 < A < \infty$ and some sequence $T = T(n) = O(A\, n^2 \log(n))$, they satisfy

$$\sup_{(X_0,Y_0)\in \mathrm{SO}(n)\times \mathrm{SO}(n)} \mathbb{P}[\|X_T - Y_T\|_{\mathrm{HS}} \le n^{-A}] \ge 1 - n^{-A}, \qquad T = O(A\, n^2 \log n).$$

The following result, obtained as a consequence of this result, will allow us to restrict our attention to starting points X_0 that satisfy $\|X_0 - Y_0\|_{\mathrm{HS}} \le \omega_n$ without loss of generality:

Lemma 2.3. *Denote by* Π *the class of measures* ν *on* $\mathrm{SO}(n) \times \mathrm{SO}(n)$ *that are supported on pairs* (x, y) *with* $\|x - y\|_{\mathrm{HS}} \le \omega_n$. *Assume that, for some* $T \in \mathbb{N}$ *and some* $0 < c < 1$,

$$\sup_{(X_0,Y_0)\sim\nu\in\Pi} \|\mathcal{L}(X_T) - \mathcal{L}(Y_T)\|_{\mathrm{TV}} \le c. \tag{2.3}$$

Then for any $A > 0$, $C > \frac{\log(\omega_n)}{\log(n)} + A$ *and* $S > \lceil n^2 \log(n) \left(\frac{1}{2} + \log(\pi) + C\right) \rceil$,

$$\sup_{X_0 = x \in \mathrm{SO}(n)} \|\mathcal{L}(X_{T+S}) - \mu\|_{\mathrm{TV}} \le c + n^{-A}. \tag{2.4}$$

Proof. This is an immediate consequence of Theorem 1 of [Oli09]. For any $X_0 = x \in \mathrm{SO}(n)$ and $Y_0 \sim \mu$, Theorem 1 of [Oli09] implies that there is a coupling of $\{(X_t, Y_t)\}_{t=0}^S$ so that

$$\mathbb{E}[\|X_S - Y_S\|_{\mathrm{HS}}^2] \le \omega_n n^{-A}.$$

By Markov's inequality, it follows that for this coupling, we can write the joint law of (X_S, Y_S) as:

$$\mathcal{L}(X_S, Y_S) = (1 - n^{-A})\nu_x + (n^{-A}) r_x \tag{2.5}$$

for some $\nu_x \in \Pi$ and some probability measure r_x on $\mathrm{SO}(n) \times \mathrm{SO}(n)$; we do not need to know the details of either of these distributions.

Now the claim follows from a standard minorization argument. Let $X_0 = x \in \mathrm{SO}(n)$, $Y_0 \sim \mu$ and Z be a Bernoulli random variable with $\mathbb{P}[Z = 0] = 1 - \mathbb{P}[Z = 1] = n^{-A}$. We couple $\{X_t\}_{t=0}^S$, $\{Y_t\}_{t=0}^S$ and Z so that

$$\mathcal{L}(X_S, Y_S) = \begin{cases} \nu_x & \text{if } Z = 1 \\ r_x & \text{if } Z = 0. \end{cases}$$

[1] The coupling referred to in Theorem 1 of [Oli09] is not the same as the coupling given in Definition 2.2, though the two are closely related.

This coupling is possible by Equation (2.5). If $Z = 1$, we then couple $\{X_t\}_{t=S}^{T+S}$, $\{Y_t\}_{t=S}^{T+S}$ so that

$$\mathbb{P}[X_{T+S} \neq Y_{T+S}] \leq c;$$

such a coupling is possible by our hypothesis stated in Equation (2.3). If $Z = 1$, we allow $\{X_t\}_{t\geq S}$, $\{Y_t\}_{t\geq S}$ to evolve independently. Under this coupling, we have

$$\begin{aligned}\mathbb{P}[X_{T+S} \neq Y_{T+S}] &= \mathbb{P}[X_{T+S} \neq Y_{T+S}|Z = 0]\mathbb{P}[Z = 0] + \mathbb{P}[X_{T+S} \neq Y_{T+S}|Z = 1]\mathbb{P}[Z = 1] \\ &\leq \mathbb{P}[Z = 0] + \mathbb{P}[X_{T+S} \neq Y_{T+S}|Z = 1] \\ &\leq n^{-A} + c.\end{aligned}$$

By Lemma 2.1, this completes the proof. □

Now we define our second stage coupling in detail. We construct our coupling by running two 'scaffolding' copies of Kac's walk (these are the chains $\{\hat{X}_t, \hat{Y}_t\}_{t\geq 0}$ defined below) according to Definition 2.2, and then insert a little bit of extra randomness at carefully chosen points (this corresponds to the choice of $\delta(x)$, $\delta(y)$ below). We first need the following definitions.

Definition 2.4 (Induced Map)**.** Fix $\mathcal{T} \in \mathbb{N}$ and a set of integers $\mathcal{S} = \{s_1 < \ldots < s_N\}$ with $0 \leq s_1 < s_N < \mathcal{T}$. Fix a sequence $\mathcal{I} = \{i_0, \ldots, i_{\mathcal{T}-1}\}$ with $i_k \in \{1, 2, \cdots, N\}$ and a sequence $\eta = \{\eta_0, \ldots, \eta_{\mathcal{T}-1}\}$ with $\eta_k \in [0, 2\pi)$. Fix $X \in \mathrm{SO}(n)$ and let $\mathcal{A} = \{X, \mathcal{T}, \mathcal{S}, \mathcal{I}, \eta\}$.

Define the maps $e_{\mathcal{A},t} : [0, 2\pi)^N \to [0, 2\pi)$ by

$$\begin{aligned}e_{\mathcal{A},t}(x_1, \ldots, x_N) &= \eta_t, & t \notin \mathcal{S} \\ e_{\mathcal{A},t}(x_1, \ldots, x_N) &= \eta_t + x_\ell, & t = t_\ell \in \mathcal{S}.\end{aligned}$$

Finally, fix $0 < c < \pi$ and define the map $f_{\mathcal{A},c} : [-c, c]^N \to \mathrm{SO}(n)$ by

$$f_{\mathcal{A},c}(x_1, \ldots, x_N) = \prod_{t=\mathcal{T}-1}^{0} \mathrm{R}(i_t, e_{\mathcal{A},t}(x_1, \ldots, x_N))X. \tag{2.6}$$

Our coupling will involve the choice of certain marked times $s_1 < s_2 < \ldots < s_N$. The following 'greedy' strategy for choosing these marked times is difficult to analyze, but gives a useful algorithm for obtaining mixing bounds via simulation (see Remark 5.3):

Definition 2.5 (Greedy Subset Choice)**.** Let $\{i(t)\}_{t\geq 0}$ be a sequence with $i(t) \in \{1, 2, \ldots, N\}$. For $t \geq 0$, define $V_t = \mathrm{span}(\{a_{i(1)}, a_{i(2)}, \ldots, a_{i(t)}\})$. Then define $s_1 = 0$ and inductively define

$$s_{\ell+1} = \inf\{s > s_\ell \,:\, V_s \neq V_{s_\ell}\}.$$

Let $T = s_N + 1$ and $\mathcal{S} = \{s_1, \ldots, s_N\}$. Note that the quantities $\mathcal{S}$ and T are both functions only of the sequence $\{i_t\}_{t\geq 0}$.

A critical step in the analysis of our Markov chain is a bound on the smallest singular value of a matrix D with highly dependent entries that is associated with our subset choice (see Equation (6.15)). Since D has highly dependent entries, we are only be able to bound its smallest singular value by comparing D to a 'limiting' random matrix D_∞ for which we can make exact computations (see Section 6 and Remark 6.7). To make this comparison possible, D must be close to D_∞. The following 'lazy' choice of marked times ensures that this happens:

Definition 2.6 (Lazy Subset Choice)**.** Fix a constant $\mathcal{Q} > 0$. Let $\{i(t)\}_{t\geq 0}$ be a sequence with $i(t) \in \{1, 2, \ldots, N\}$. We inductively define the sequence $\{s_\ell\}_{\ell=1}^N$ by setting

$$s_1 = \min\{t \geq 0 \,:\, i(t) = 1\}$$
$$s_{\ell+1} = \min\{t \geq s_\ell + \mathcal{Q}n^2 \log(n) \,:\, i(t) = \ell + 1\}.$$

Define

$$\mathcal{S} = \{s_1, \ldots, s_N\} \tag{2.7}$$

and set $T = s_N + 1$.

2.1. **Non-Markovian Coupling.** We now give a non-Markovian coupling for two copies $\{X_t\}_{t\geq 0}$, $\{Y_t\}_{t\geq 0}$ of Kac's walk. Although this coupling will be used in the proof of Theorem 1.1 for two chains that are started at nearby points, the coupling itself makes sense for any pair of starting points (X_0, Y_0). The coupling depends on the parameters $\mathcal{Q} \in \mathbb{R}^+$ and $\epsilon_n \in (0, 1)$.

As with the coupling in Definition 2.2, we define this coupling in terms of the associated update sequences $\{i_t(x), \theta_t(x)\}_{t\geq 0}$ and $\{i_t(y), \theta_t(y)\}_{t\geq 0}$:

(1) For $t \geq 0$, choose $i_t \sim \text{unif}\{1, 2, \ldots, N\}$ and set
$$i_t(y) = i_t(x) = i_t. \tag{2.8}$$
(2) Choose $\mathcal{S}$ and T as in Definition 2.6.
(3) We couple two copies $\{\hat{X}_t, \hat{Y}_t\}_{t\geq 0}$ of Kac's walk started at $\hat{X}_0 = X_0$, $\hat{Y}_0 = Y_0$ according to Definition 2.2 and using the same sequence $\{i_t(x), i_t(y)\}_{t\geq 0}$; let $\{\eta_t(x), \eta_t(y), h_t\}_{0\leq t<T}$ be as in Definition 2.2. These walks 'scaffold' our coupling of $\{X_t, Y_t\}_{t\geq 0}$.
(4) Define
$$\mathcal{A} = \{X_0, T, \mathcal{S}, \{i_0, \ldots, i_{T-1}\}, \{\eta_0(x), \ldots, \eta_{T-1}(x)\}\}, \tag{2.9}$$
$$\mathcal{B} = \{Y_0, T, \mathcal{S}, \{i_0, \ldots, i_{T-1}\}, \{\eta_0(y), \ldots, \eta_{T-1}(y)\}\}.$$
Let $\{\delta_t(x)\}_{1\leq t\leq N}$, $\{\delta_t(y)\}_{1\leq t\leq N}$ be two sequences of i.i.d. $\text{unif}(-\epsilon_n, \epsilon_n)$ random variables and write $\delta(z) = (\delta_1(z), \ldots, \delta_N(z))$ for $z \in \{x, y\}$. We couple the two sequences $\delta(x)$, $\delta(y)$ so that they satisfy
$$\mathbb{P}[f_{\mathcal{A},\epsilon_n}(\delta(x)) = f_{\mathcal{B},\epsilon_n}(\delta(y))] = \|\mathcal{L}(f_{\mathcal{A},\epsilon_n}(\delta(x))) - \mathcal{L}(f_{\mathcal{B},\epsilon_n}(\delta(y)))\|_{\text{TV}}. \tag{2.10}$$
For $t < T$, set
$$\begin{aligned} \theta_t(x) &= \eta_t(x), && t \notin \mathcal{S} \\ \theta_t(x) &= \eta_t(x) + \delta_\ell(x), && t = t_\ell \in \mathcal{S}, \end{aligned} \tag{2.11}$$
and similarly
$$\begin{aligned} \theta_t(y) &= \eta_t(y), && t \notin \mathcal{S} \\ \theta_t(y) &= \eta_t(y) + \delta_\ell(y), && t = t_\ell \in \mathcal{S}. \end{aligned} \tag{2.12}$$
This completes our construction of $\theta_t(x)$, $\theta_t(y)$ for $0 \leq t < T$.
(5) To construct $\theta_t(x)$, $\theta_t(y)$ for $t \geq T$, let $\{\theta_t\}_{t\geq T}$ be a sequence of i.i.d. random variables with $\theta_t \sim \text{unif}[0, 2\pi)$ and set $\theta_t(x) = \theta_t(y) = \theta_t$.

Thus Equation (2.8) couples $i_t(x)$ and $i_t(y)$ and Equations (2.11) and (2.12) give the coupling of $\theta_t(x)$ and $\theta_t(y)$. The stochastic processes $\{X_t\}$ and $\{Y_t\}$ are coupled through this non-Markovian coupling of update sequences.

Lemma 2.7 (Validity of Coupling)**.** *The stochastic process* $\{X_t\}_{t\geq 0}$, $\{Y_t\}_{t\geq 0}$ *have the same distribution as Kac's walk.*

Proof. We give the proof for $\{X_t\}_{t\geq 0}$, as the proof for $\{Y_t\}_{t\geq 0}$ is identical. To check that $\{X_t\}_{t\geq 0}$ has the same distribution as Kac's walk, it is sufficient to check that the update sequence $\{(i_t(x), \theta_t(x))\}_{t\geq 0}$ has the correct distribution. The sequence $\{i_t(x)\}_{t\geq 0}$ is defined to be an i.i.d. sequence of unif$(\{1, 2, \ldots, N\})$ random variables, so this is correct.

To check that $\{\theta_t(x)\}_{t\geq 0}$ is an i.i.d. sequence of unif$[0, 2\pi)$ random variables conditional on $\{i_t(x)\}_{t\geq 0}$, we note that conditional on $\{\delta_t(x)\}_{t=1}^N$ and $\{i_t(x)\}_{t\geq 0}$, $\{\eta_t(x)\}_{t\geq 0}$ is an i.i.d. sequence of unif$[0, 2\pi)$ random variables. Since any random variable plus an independent unif$[0, 2\pi)$ random variable is itself uniform on $[0, 2\pi)$, this implies $\{\theta_t(x)\}_{t\geq 0}$ is an i.i.d. sequence of unif$[0, 2\pi)$ random variables. □

Remark 2.8. The non-Markovian coupling in Section 2.1 makes sense for any Gibbs sampler. To use it, one only requires

(1) a representation of the Gibbs sampler in terms of the sequence of randomly-chosen coordinates to be updated $\{i_t\}_{t\geq 0}$ and the random quantile $\{\theta_t\}_{t\geq 0}$, and
(2) a candidate contractive coupling to form the 'scaffold.'

These requirements are not onerous; (1) is the usual way to write down a Gibbs sampler, while any optimal 1-step coupling is a reasonable candidate for (2).

2.2. **Informal Description of Coupling Analysis.** By construction, we have $X_T = f_{\mathcal{A},\epsilon_n}(\delta(x))$, $Y_T = f_{\mathcal{B},\epsilon_n}(\delta(y))$. Consequently, any bound on the total variation distance in Equation (2.10) immediately gives a bound on the coupling probability required by Lemma 2.1. To prove a mixing bound, it is thus enough to prove that the images $I_X = f_{\mathcal{A},\epsilon_n}([-\epsilon_n, \epsilon_n]^N)$ and $I_Y = f_{\mathcal{B},\epsilon_n}([-\epsilon_n, \epsilon_n]^N)$ of $f_{\mathcal{A},\epsilon_n}$ and $f_{\mathcal{B},\epsilon_n}$ have a large overlap and the respective Jacobians $M_{\mathcal{A}}$ and $M_{\mathcal{B}}$ are roughly constant. By Taylor expansion, these images are roughly of the form

$$I_X \approx f_{\mathcal{A},\epsilon_n}(0) + \exp(M_{\mathcal{A}}[-\epsilon_n, \epsilon_n]^N),\ I_Y \approx f_{\mathcal{B},\epsilon_n}(0) + \exp(M_{\mathcal{B}}[-\epsilon_n, \epsilon_n]^N).$$

By the assumption that $\|X_0 - Y_0\| \leq \omega_n$ is small and the main result of [Oli09], we have $\|f_{\mathcal{A},\epsilon_n}(0) - f_{\mathcal{B},\epsilon_n}\| = O(\omega_n)$ and $\|M_{\mathcal{A}} - M_{\mathcal{B}}\| = O(\omega_n)$. Thus, to check that $I_X \approx I_Y$, it is enough to check that the smallest singular value $\sigma_1(M_{\mathcal{A}})$ of $M_{\mathcal{A}}$ satisfies $\sigma_1(M_{\mathcal{A}}) \gg \|M_{\mathcal{A}} - M_{\mathcal{B}}\|$. Similar bounds show that the Jacobians are nearly constants. This heuristic is formalized in Section 4.

3. Technical Lemmas

We give a collection of general estimates that will be used throughout the paper. The proofs are deferred to Appendix A, and can be safely skipped on a first reading of this paper.

3.1. **Matrix Estimates.** We use the following result repeatedly.

Lemma 3.1. *Fix* $k \geq 1$ *and let* $P_1, \ldots, P_k$ *and* $Q_1, \ldots, Q_k$ *be two sequences of elements of* $M(n)$*. Then*

$$\Big\| \prod_{i=1}^k Q_i - \prod_{i=1}^k P_i \Big\|_{\mathrm{HS}} \leq \sum_{i=1}^k \Big\| \prod_{\ell=1}^{i-1} Q_\ell \Big\|_{\mathrm{Op}} \|Q_i - P_i\|_{\mathrm{HS}} \Big\| \prod_{\ell=i+1}^k P_\ell \Big\|_{\mathrm{Op}}. \tag{3.1}$$

We also have the elementary bound:

Lemma 3.2 (Determinant Estimate)**.** *Let M_1, M_2 be two N by N symmetric matrices. For a general matrix M, denote by $\sigma_1(M) \leq \sigma_2(M) \leq \ldots \leq \sigma_N(M)$ the ordered singular values of M. Assume that*

$$\|M_1 - M_2\|_{\mathrm{Op}} \leq \delta \sigma_1(M_1) \tag{3.2}$$

for some $0 < \delta < 1$. Then the determinants of M_1, M_2 satisfy:

$$|\frac{\det(M_2)}{\det(M_1)} - 1| \leq N^{\frac{N}{2}} \delta^N.$$

Our next collection of bounds will require further notation. Define two functions $f : [-c,c]^N \mapsto \mathrm{SO}(n)$ and $g : [-c,c]^N \mapsto \mathrm{SO}(n)$:

$$f(x_1, x_2, \cdots, x_N) = \prod_{k=1}^{N} \mathrm{R}_k \exp\big((\theta_k + x_k)a_k\big), \tag{3.3}$$

$$g(y_1, y_2, \cdots, y_N) = \prod_{k=1}^{N} \tilde{\mathrm{R}}_k \exp\big((\tilde{\theta}_k + y_k)a_k\big),$$

where $\theta_k, \tilde{\theta}_k \in [0, 2\pi)$, $\mathrm{R}_k, \tilde{\mathrm{R}}_k \in \mathrm{SO}(n)$ and $a_k \in \mathfrak{so}(n)$. For $1 \leq i+1 < j \leq N$, define

$$M_{i,j} = M_{i,j}(x_1, \ldots, x_N) \equiv \prod_{k=i+1}^{j-1} \mathrm{R}_k \, e^{(\theta_k + x_k)a_k},$$

for $1 \leq j+1 < i \leq N$, define

$$M_{i,j} = \prod_{k=j+1}^{i-1} R_k \, e^{(\theta_k + x_k)a_k},$$

and for $1 \leq i < N$ define $M_{i,i+1} = M_{i+1,i} = \mathrm{Id}$. The derivative map $\mathrm{d}f : T_p[-c,c]^N \mapsto T_{f(p)}\mathrm{SO}(n)$ in the direction $h = \sum_k h_k \frac{\partial}{\partial e_k}$ is

$$\mathrm{d}f_x(h) = \sum_{j=1}^{N} \prod_{k=1}^{N} \mathrm{R}_k \, e^{(\theta_k + x_k)a_k} (\mathrm{Id} + (h_k a_k - \mathrm{Id}) 1_{k=j}). \tag{3.4}$$

We state two lemmas that bound various approximations related to f and g:

Lemma 3.3 (Closeness of Tangent Maps)**.** *Let $f : [-c,c]^N \mapsto \mathrm{SO}(n)$ be as defined in Equation (3.3). Assume that $c < N^{-3}$ and $\max_k |\theta_k - \tilde{\theta}_k| \leq c^2$, where $\theta_k, \tilde{\theta}_k \in [0, 2\pi)$ are as in Equation (3.3). Then for all $x, y \in [-c,c]^N$ and all $h \in \mathbb{R}^N$ with $\|h\| \leq 1$,*

$$\|\mathrm{d}f_x(h) - \mathrm{d}f_y(h)\|_{\mathrm{HS}} \leq 4N^2 c \tag{3.5}$$

$$\|\mathrm{dL}_{f(y)(f(x))^{-1}} \mathrm{d}f_x(h) - \mathrm{d}f_y(h)\|_{\mathrm{HS}} \leq 8N^2 c.$$

Lemma 3.4 (Approximation by Exponential Map)**.** *Let f be of the form given in Equation (3.3). Then for $c < N^{-3}$,*

$$\|f(x) - f(0)\exp(\mathrm{dL}_{f(0)^{-1}} \mathrm{d}f_0(x))\|_{\mathrm{HS}}^2 \leq 8N^2 c^2.$$

We need the following expression for the Jacobian of f and g:

Lemma 3.5 (Jacobian Formula)**.** *Let f be the function defined in Equation (3.3). For $1 \le i < j \le N$, define*

$$D_{i,j} = -\mathrm{Tr}[a_i M_{i,j} \mathrm{R}_j a_j \mathrm{R}_j^{-1} M_{i,j}^{-1}]; \tag{3.6}$$

for $1 \le j < i \le N$, define

$$D_{i,j} = D_{j,i}.$$

Finally, set $D_{i,i} = 1$ for all $1 \le i \le N$ and let D be the matrix with entries $[D_{i,j}]$. Then for all $h = (h_1, \ldots, h_N) \in \mathbb{R}^N \backslash \{0\}$ and all $x \in [-c, c]^N$,

$$\langle \mathrm{d}f_x(h), \mathrm{d}f_x(h') \rangle_{\mathrm{HS}} = \langle h, h' \rangle + \sum_{i \ne j} h_i h'_j D_{i,j} = h^\dagger D h'. \tag{3.7}$$

We point out that f is generally a diffeomorphism for c sufficiently small:

Lemma 3.6. *Let f be of the form given in Equation (3.3) and assume that it satisfies*

$$\begin{aligned} \inf_{h \ne 0} \frac{2\|df_x(h)\|}{\|h\|} &\ge \phi_n \\ \sup_{h \ne 0} \frac{\|df_x(h)\|}{\|h\|} &\le N, \end{aligned} \tag{3.8}$$

for some sequence

$$\phi_n \le (2n)^{-30}. \tag{3.9}$$

Then there exist constants $0 < C_0 < 1$, $N_0 < \infty$ that do not depend on n or the particular sequence $\{\phi_n\}_{n \ge 0}$ satisfying (3.9), such that the function f is a diffeomorphism whenever $c < C_0 n^{-6} \phi_n$ for all $n > N_0$ sufficiently large.

3.2. **Probability Estimates.** We check that the timescale T of the non-Markovian coupling defined in Section 2.1 is not too large:

Lemma 3.7. *Let s_N be defined as in Definition 2.5. For all $c \in \mathbb{R}$,*

$$\lim_{n \to \infty} \mathbb{P}[s_N < N \log(N) + cN] = e^{-e^c}. \tag{3.10}$$

Lemma 3.8. *Let s_N be defined as in Definition 2.6. Then for all $k \ge 2$,*

$$\mathbb{P}[s_N > \mathcal{Q} N^2 \lceil \log(N) \rceil + kN^2] \le e^{-\frac{N}{4}}. \tag{3.11}$$

We check that the chains in Section 2.1 will stay very close to each other, with high probability:

Lemma 3.9 (Closeness of Paths)**.** *Assume that X_0, Y_0 satisfy $\|X_0 - Y_0\|_{\mathrm{HS}} \le \omega_n$ for some $0 < \omega_n \le n^{-2700}$. Let $\{X_t, Y_t, \hat{X}_t, \hat{Y}_t\}_{t \ge 0}$ be coupled as in Definition 2.1, with $c \equiv \epsilon_n \in [\omega_n^{\frac{1}{30}}, n^{-900}]$. Fix $0 < C < \frac{\log(\omega_n)}{\log(n)} - 4$. Then*

$$\mathbb{P}[\sup_{0 \le t \le n^5} \|\hat{X}_t - \hat{Y}_t\|_{\mathrm{HS}} \le 2\|\hat{X}_0 - \hat{Y}_0\|_{\mathrm{HS}} n^{5+C}] \ge 1 - n^{-C} \tag{3.12}$$

$$\sup_{0 \le t \le n^5} \|X_t - \hat{X}_t\|_{\mathrm{HS}} \le 6n^5 \epsilon_n \tag{3.13}$$

$$\sup_{0 \le t \le n^5} \|Y_t - \hat{Y}_t\|_{\mathrm{HS}} \le 6n^5 \epsilon_n. \tag{3.14}$$

4. Coupling Argument

We now give a generic bound on the total variation distance in Equation (2.10), under the assumption that nothing 'too bad' happened. This formalizes the argument in Section 2.2.

Lemma 4.1 (Coupling Construction)**.** *Fix sequences $\phi_n \leq (2n)^{-30}$, $\epsilon_n = \phi_n^{30}$ and $\omega_n = \epsilon_n^{30}$. Fix a constant C that satisfies $0 < C < \frac{\log(\omega_n)}{\log(\epsilon_n)} - 6$ for all $n \in \mathbb{N}$, and let $f = f_{\mathcal{A},\epsilon_n}$ and $g = f_{\mathcal{B},\epsilon_n}$ be as defined in Equation (2.6), with $\mathcal{A}$, $\mathcal{B}$ be as in equation (2.9). Assume that*

$$\|f(0) - g(0)\|_{\mathrm{HS}} \leq \omega_n n^{5+C} \tag{4.1}$$

and that for $q \in \{f, g\}$, we have:

$$\begin{aligned} \inf_{h \neq 0} \frac{2\|dq_x(h)\|}{\|h\|} &\geq \phi_n \\ \sup_{h \neq 0} \frac{\|dq_x(h)\|}{\|h\|} &\leq N. \end{aligned} \tag{4.2}$$

Then it is possible to couple two sequences of i.i.d. random variables $U_1, \ldots, U_N \sim \mathrm{unif}(-\epsilon_n, \epsilon_n)$ and $V_1, \ldots, V_N \sim \mathrm{unif}(-\epsilon_n, \epsilon_n)$ so that

$$\mathbb{P}[f(U_1, \ldots, U_N) \neq g(V_1, \ldots, V_N)] \leq 513 N^2 \phi_n^{-1} \epsilon_n = o(1).$$

Proof. By Lemma 3.4,

$$\begin{aligned} \|f(x) - f(0)\exp(\mathrm{dL}_{f(0)^{-1}} \circ \mathrm{d}f_0(x))\|_{\mathrm{HS}} &\leq 8N^2\epsilon_N^2 \\ \|g(x) - g(0)\exp(\mathrm{dL}_{g(0)^{-1}} \circ \mathrm{d}g_0(x))\|_{\mathrm{HS}} &\leq 8N^2\epsilon_N^2. \end{aligned} \tag{4.3}$$

for all $x \in (-\epsilon_n, \epsilon_n)^N$. For $0 < r < \infty$, define

$$\begin{aligned} \mathcal{H}_f(r) &= \{f(0)\exp(\mathrm{dL}_{f(0)^{-1}} \circ \mathrm{d}f_0(x)) \,:\, x \in [-r, r]^N\} \\ \mathcal{H}_g(r) &= \{g(0)\exp(\mathrm{dL}_{g(0)^{-1}} \circ \mathrm{d}g_0(x)) \,:\, x \in [-r, r]^N\}. \end{aligned}$$

We claim:

Proposition 4.2. *For $c > 0$, set*

$$\begin{aligned} u_1 &= c - 32N^2\phi_n^{-1}\epsilon_n^2 \\ u_2 &= c + 32N^2\phi_n^{-1}\epsilon_n^2. \end{aligned} \tag{4.4}$$

Then for all n sufficiently large,

$$\begin{aligned} \mathcal{H}_f(u_1) \subset f([-c, c]^N) \subset \mathcal{H}_f(u_2) \\ \mathcal{H}_g(u_1) \subset g([-c, c]^N) \subset \mathcal{H}_g(u_2), \end{aligned} \tag{4.5}$$

uniformly in sequences $c = c_n$ satisfying $\frac{1}{2}\epsilon_n \leq c \leq 2\epsilon_n$.

Proof. Assume without loss of generality that $f(0) = \mathrm{Id}$. Fix $x \in \partial[-c, c]^N$. Since the exponential map is surjective (see, *e.g.*, Theorem 6.9.3 of [AM]), we can write $f(x) = \exp(h)$. By taking a solution h to $f(x) = \exp(h)$ with small norm, we can also assume that $\|h\|_{\mathrm{HS}} \leq 2N^2\phi_n^{-1}\epsilon_n$. Since $\mathrm{d}f_0$ has full rank, we can write $h = \mathrm{d}f_0(y)$ for some y. We calculate:

$$\begin{aligned} \|x - y\|_\infty \leq \|x - y\| &\leq 2\phi_n^{-1}\|\mathrm{d}f_0(x) - \mathrm{d}f_0(y)\|_{\mathrm{HS}} \\ &= 2\phi_n^{-1}\|h - \mathrm{d}f_0(x)\|_{\mathrm{HS}} \end{aligned} \tag{4.6}$$

$$\begin{aligned}
&\leq 4\phi_n^{-1}\|\exp(h) - \exp(\mathrm{d}f_0(x))\|_{\mathrm{HS}} \\
&= 4\phi_n^{-1}\|f(x) - \exp(\mathrm{d}f_0(x))\|_{\mathrm{HS}} \\
&\leq 32N^2\phi_n^{-1}\epsilon_n^2.
\end{aligned}$$

The first line follows from inequality (4.2). The third line follows for n sufficiently large from a Taylor expansion and the fact that $\epsilon_n \ll n^{-4}$. The last line follows from inequality (4.3). This proves that if $x \in \partial[-c,c]^N$ and $f(x) = \exp(\mathrm{d}f_0(y))$, then $y \notin [-u_1,u_1]^N$.

This implies that $f(\partial[-c,c]^N) \subset \mathcal{H}_f(u_1)^c$. Since f is a diffeomorphism by Lemma 3.6 and a map between manifolds of the same dimension N, this implies $f(\partial[-c,c]^N) = \partial f([-c,c]^N) \subset \mathcal{H}_f(u_1)^c$. Using the fact that both f and the exponential map are diffeomorphisms, the condition $\partial f([-c,c]^N) \subset \mathcal{H}_f(u_1)^c$ together with the fact that $f(0) \subset f([-c,c]^N) \cap \mathcal{H}_f(u_1) \neq \emptyset$, implies the containment condition $\mathcal{H}_f(u_1) \subset f([-c,c]^N)$. This is exactly the left-hand side of the first containment condition (4.5).

To prove the right-hand side (4.5), essentially the same calculation shows that for any $p = \exp(\mathrm{d}f_0(y)) \in f([-c,c]^N)$, we have $y \in [-u_2,u_2]^N$. This immediately implies the right-hand side of the first containment condition (4.5).

The proof of the second containment condition (4.5) is identical, so this completes the proof of the proposition. □

Since the exponential map is surjective and $\mathrm{d}f_0$ has full rank, there exists some h so that $g(0) = \mathrm{L}_{f(0)}\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(h))$. By the fact that the exponential map takes geodesic paths and the assumption that $\epsilon_n = o(n^{-5})$, followed by inequality (4.1), we have

$$\|h\|_{\mathrm{HS}} \leq \|f(0) - g(0)\|_{\mathrm{HS}} \leq \omega_n n^{5+C}. \tag{4.7}$$

We claim:

Proposition 4.3. *Fix $c > 0$, set u_1, u_2 as in equation (4.4), and set*

$$\begin{aligned}
v_1 &= u_1 - 64N^2\phi_n^{-1}\epsilon_n^2 \\
v_2 &= u_2 + 64N^2\phi_n^{-1}\epsilon_n^2.
\end{aligned}$$

Then for all n sufficiently large,

$$\mathcal{H}_f(v_1) \subset \mathcal{H}_g(u_1) \subset \mathcal{H}_g(u_2) \subset \mathcal{H}_f(v_2) \tag{4.8}$$

uniformly in $c = c_n$ satisfying $\frac{1}{2}\epsilon_n \leq c \leq 2\epsilon_n$.

Proof. Assume that the left-most containment in (4.8) is not true. Then there exists $p \in \mathcal{H}_f(v_1)\backslash\mathcal{H}_g(u_1)$. Write $p = f(0)\exp(\mathrm{dL}_{f(0)^{-1}} \circ \mathrm{d}f_0(x))$ for some $x \in [-v_1,v_1]^N$. By the local surjectivity of the exponential map, we can write $p = g(0)\exp(\mathrm{dL}_{g(0)^{-1}} \circ \mathrm{d}g_0(y))$ for some y. Since $p \notin \mathcal{H}_g(u_1)$, we have $y \notin [-u_1,u_1]^N$. However, by essentially the same calculation as in (4.6) combined with inequality (4.7),

$$\|x - y\|_\infty \leq \|x - y\| \leq 32N^2\phi_n^{-1}\epsilon_n^2 + \omega_n n^{5+C} \leq 64N^2\phi_n^{-1}\epsilon_n^2.$$

Thus,

$$v_1 \geq \|x\|_\infty \geq \|y\|_\infty - 64N^2\phi_n^{-1}\epsilon_n^2 > u_1 - 64N^2\phi_n^{-1}\epsilon_n^2 = v_1.$$

This is a contradiction, and so no such p exists. This completes the proof of the first containment relationship in (4.8). The second containment is trivial, and the third is proved in essentially the same way as the first. This completes the proof of the proposition. □

Combining Propositions 4.2 and 4.3, we have

$$f([-\epsilon_n, \epsilon_n]^N) \supset \mathcal{H}_g(c_1) \supset g([-c_2, c_2]^N), \tag{4.9}$$

where $c_1 = \epsilon_n - 256N^2\phi_n^{-1}\epsilon_n^2$ and $c_2 = c_1 - 256N^2\phi_n^{-1}\epsilon_n^2$.

Let ρ_f and ρ_g denote the densities of $\mathcal{L}(f(U_1, \ldots, U_N))$ and ρ_g of $\mathcal{L}(g(V_1, \ldots, V_N))$. By Lemma 3.2,

$$|\frac{\rho_g(x)}{\rho_f(y)} - 1| \leq N^{\frac{N}{2}} \left(\frac{\|df_x - dg_y\|_{\mathrm{Op}}}{\sigma_1(df_x)} \right)^N.$$

By Lemma 3.3 and assumption (4.2), $\frac{\|df_x - dg_y\|_{\mathrm{Op}}}{\sigma_1(df_x)} \leq 32\phi_n^{-1}N^2\epsilon_n$, and so

$$|\frac{\rho_g(x)}{\rho_f(y)} - 1| \leq (32N^{2.5}\phi_n^{-1}\epsilon_n)^N \ll N^{-N}. \tag{4.10}$$

Combining inequality (4.10) with the containment condition (4.9),

$$\begin{aligned}
\|\mathcal{L}(f(U_1, \ldots, U_N)) - \mathcal{L}(g(V_1, \ldots, V_N))\|_{\mathrm{TV}} &\leq 1 - \frac{\mu(f([-\epsilon_n, \epsilon_n]^N) \cap g([-\epsilon_n, \epsilon_n]^N))}{\mu(g([-\epsilon_n, \epsilon_n]^N))}(1 + N^{-N}) \\
&\leq 1 - \frac{\mu(g([-c_2, c_2]^N))}{\mu(g([-\epsilon_n, \epsilon_n]^N))}(1 + N^{-N}). \qquad (4.11) \\
&= 1 - \frac{(2c_2)^N}{(2\epsilon_N)^N}(1 + N^{-N}) \\
&\leq 1 - \frac{(2\epsilon_n - 1024N^2\phi_n^{-1}\epsilon_n^2)^N}{(2\epsilon_n)^N}(1 + N^{-N}) \\
&\leq 1 - 513N^2\phi_n^{-1}\epsilon_n.
\end{aligned}$$

This completes the proof of the lemma. □

5. Relating Mixing Times to Singular Values

Let $f_{\mathcal{A},\epsilon_n}$ be as in Equation 2.6, and let D be the Jacobian of $f_{\mathcal{A},\epsilon_n}$ at $(0, 0, \ldots, 0)$. In this section, we apply our current bound to relate the mixing time of Kac's walk to the smallest singular value $\sigma_1(D)$ of the matrix D. Define the scaling sequence $\{\phi_n\}_{n\geq 1}$ by

$$\phi_n = \min((2n)^{-30}, \inf\{r > 0 : \mathbb{P}[\sigma_1(D) \leq r] < \frac{1}{\sqrt{n}}). \tag{5.1}$$

We give most of the proof of Theorem 1.1, deferring the proof of a bound on ϕ_n to Section 6. We begin with the lower bound:

Theorem 1 (Lower Bound). *For $T < N$ and any $X_0 = x \in \mathrm{SO}(n)$,*

$$\|\mathcal{L}(X_T) - \mu\|_{\mathrm{TV}} = 1.$$

Proof. The proof is essentially a matter of counting dimensions.

Let $\{X_t\}_{t\geq 0}$ be a copy of Kac's walk, and let $\{i_t, \theta_t\}_{t\geq 0}$ be the associated update variables as described in Equation (1.1). For $t \in \mathbb{N}$, define $f_t : [0, 2\pi)^t \mapsto \mathrm{SO}(n)$ by

$$f_t(x_0, \ldots, x_{t-1}) = \prod_{s=t-1}^{0} \mathrm{R}(i_s, x_s)X_0$$

and define the set $A_t = A_t(\{i_t\}_{t\geq 0}) \equiv f_t([0,2\pi)^t)$. We note that $X_t \in A_t$. From this definition, $\mathrm{Rank}((\mathrm{d}f_t)_p) \leq t < N$ for all p. By Sard's theorem (see pp. 205 of [GP74]), this implies

$$\pi(A_t) = 0. \tag{5.2}$$

Next, define

$$\mathrm{Max}_T = \cup_{I\in\{1,2,\ldots,N\}^T} A_T(I).$$

Since this is a union of only $\binom{n}{2}^T < \infty$ elements, Equation(5.2) implies

$$\pi(\mathrm{Max}_T) = 0.$$

Thus,

$$\|\mathcal{L}(X_T) - \pi\|_{\mathrm{TV}} \geq |\mathbb{P}[X_t \in \mathrm{Max}_T] - \pi(\mathrm{Max}_T)| = 1,$$

finishing the proof. □

Remark 5.1 (Dimension Counting and Curved Spaces)**.** A natural approach for obtaining a lower bound is to count the dimension of the tangent map associated with the function f_t at 0 rather than bounding the dimension of the image of f_t itself. This approach suggests that the chain will not have mixed until the first time T that the span of $\{a_{i_1}, a_{i_2}, \ldots, a_{i_T}\}$ has dimension N with high probability; by the usual coupon collector argument, this requires $T \gtrsim n^2\log(n)$.

While this approach works for Gibbs samplers on convex sets (for which the dimension of the tangent map of f_t at 0 is an upper bound on the dimension of the image of f_t), and it works for Kac's walk on the sphere for different reasons (see [PS15]), it does *not* work for Kac's walk on the sphere. In particular, it is possible for f_t to have full dimension N, despite $\mathrm{Dimension}(\mathrm{span}(\{a_{i_1}, a_{i_2}, \ldots, a_{i_t}\})) < N$. See the famous 'Euler angle' decomposition of SO(3) for an illustration of this fact [GS01].

The upper bound in Theorem 1.1 will be a corollary of the following result:

Theorem 2 (Intermediate Bound on the Mixing Time of Kac's Walk on SO(n))**.** *Fix $0 < \mathcal{Q} < \infty$ and consider the coupling given in Section 2.1. Let ϕ_n be as defined in Equation (5.1), and let $R = 8\mathcal{Q}n^4\log(n)$. Assume that $\phi_n = o(n^{n^{2.2}})$. Then*

$$\limsup_{n\to\infty} \sup_{X_0 = x\in \mathrm{SO}(n)} \|\mathcal{L}(X_t) - \mu\|_{\mathrm{TV}} = 0 \tag{5.3}$$

for any sequence $t = t(n) > R + 5n^2\log(n) + 900n^2\log(\phi_n)$.

Proof. Let ϕ_n be as in Equation (5.1), let $\epsilon_n = \phi_n^{30}$ and let $\omega_n = \epsilon_n^{30}$. Fix $R < S < N^{4.9}$. We couple two copies of Kac's walk $\{X_t\}_{t=0}^S$, $\{Y_t\}_{t=0}^S$ started at positions satisfying $\|X_0 - Y_0\|_{\mathrm{HS}} \leq \omega_n n^{-5}$ according to the non-Markovian coupling given in Section 2.1.

Let $\mathcal{E}_1$ be the event that $f_{\mathcal{A},\epsilon_n}$ and $f_{\mathcal{B},\epsilon_n}$ are well-defined at time S and condition (4.1) is satisfied at time S with constant $C = \frac{\log(\omega_n)}{\log(\epsilon_n)} - 10 = 20$, and let $\mathcal{E}_2$ be the event that $f_{\mathcal{A},\epsilon_n}$ and $f_{\mathcal{B},\epsilon_n}$ are well-defined at time S and condition (4.2) is satisfied at time S. The conclusion of Lemma 4.1, together with Remark 2.2, implies that

$$\|X_S - Y_S\|_{\mathrm{TV}} \leq (1 - \mathbb{P}[\mathcal{E}_1] - \mathbb{P}[\mathcal{E}_2]) + o(1). \tag{5.4}$$

By Lemma 3.8, $\mathbb{P}[s_N < R] = 1 - o(1)$, and so the functions $f_{\mathcal{A},\epsilon_n}$ and $f_{\mathcal{B},\epsilon_n}$ are well-defined for all times $S > R$ with probability $1 - o(1)$. By Lemma 3.9, the condition (4.1) of Lemma 4.1 is satisfied with probability $1 - o(1)$. Thus, $\mathbb{P}[\mathcal{E}_1] = 1 - o(1)$. By the definition of ϕ_n, we have that $\mathbb{P}[\mathcal{E}_2] = 1 - o(1)$. Combining these bounds with Inequality (5.4), we conclude that

$$\|X_S - Y_S\|_{\mathrm{TV}} = o(1).$$

By Lemma 2.3, this implies for $\{X_t\}_{t\geq 0}$ a copy of Kac's walk starting at any point $X_0 = x \in \mathrm{SO}(n)$ and $T > R + 5n^2 \log(n) + 900n^2 \log(\phi_n)$,

$$\sup_{X_0 = x \in \mathrm{SO}(n)} \|\mathcal{L}(X_T) - \mu\|_{\mathrm{TV}} = o(1).$$

This completes the proof. □

The following, similar, result allows us to calculate much sharper numerical bounds on the mixing time of Kac's walk:

Theorem 3 (Alternative Intermediate Bound on the Mixing Time of Kac's Walk on $\mathrm{SO}(n)$)**.** *Consider the coupling given in Section 2.1, with the reference to Definition 2.6 replaced by a reference to Definition 2.5. Let ϕ_n be as defined in Equation (5.1). Then*

$$\limsup_{n\to\infty} \sup_{X_0 = x \in \mathrm{SO}(n)} \|\mathcal{L}(X_t) - \mu\|_{\mathrm{TV}} = 0 \tag{5.5}$$

for any sequence $t = t(n)$ satisfying $\lim_{n\to\infty} \frac{t(n)}{n^2 \log(n\phi_n)} = \infty$.

Proof. The proof is almost identical to that of Theorem 2. The main difference is that all references to Lemma 3.8 should be replaced by references to Lemma 3.7.

□

Remark 5.2. In either case, we note that this implies the mixing time of Kac's walk satisfies $\tau_{\mathrm{mix}} = O(n^2 \log(\phi_n))$.

Remark 5.3 (Estimating Mixing Times via Simulation)**.** We do not analyze the bound from Theorem 3 in this paper. We include it because it can be used to obtain rigorous upper bounds on the mixing time of Kac's walk by simulation. In particular, the random matrix D that appears in Equation (5.1) can easily be simulated on a computer. Thus, the quantiles of the distribution of $\sigma_1(D)$ can be estimated by simulation, which allows us to calculate upper bounds on the constant ϕ_n with high confidence.

We point out that the existence of a method to rigorously bound the mixing time of a Gibbs sampler is not obvious. Indeed, the authors are not aware of *any* finite computation that allows one to rigorously bound the mixing time of a generic Gibbs sampler on a continuous non-convex state space. We also mention that, just as the coupling in Section 2.1 makes sense for any Gibbs sampler, this computational approach to bounding the mixing time can be extended (with some effort) to many other Gibbs samplers.

6. Singular Values of Random Matrices

The last ingredient in the proof of our upper bound on the mixing time of Kac's walk is a *lower* bound on the smallest singular value $\sigma_1(D)$ of the matrix D defined in Equation (6.15). In this section, we obtain the required bound.

We begin by giving a generic bound on the smallest singular value of a random matrix whose entries have continuous but strongly dependant entries (see Section 6.1), then apply this bound to a simple random matrix D_∞ for which some exact calculations are possible (see Section 6.2), and finally compare the smallest singular value of D_∞ to that of the matrix D defined in Equation (6.15) (see Section 6.3). This argument gives us a lower bound on the constant ϕ_n that is defined in Equation (5.1), and which has a critical role in the conclusion of Theorem 2.

We believe that the bounds in this section may be of independent interest. The notation used in this section is also essentially independent of the notation of the remainder of the paper, except where explicitly noted. Our main abstract results, given in Lemmas 6.3 and 6.4, are qualitatively similar to the main bounds in [FG13]. While our assumptions are similar to those in [FG13], the key difference is that our results apply to matrices with a great deal more dependence and which may be symmetric. Finally, the paper [FV15] gives related but much stronger conclusions than our paper or [FG13], but under much stronger independence assumptions.

6.1. **Bounds on Determinant of Random Matrices.** Let M be an $n \times n$ symmetric random matrix with associated σ-algebra (Ω, Σ). Let $\mathcal{F}_0$ denote the σ-algebra generated by the entries of M. For $1 \leq i \leq n-1$, let $\mathcal{F}_i$ be a σ-algebra under which $M[k,\ell]$ is $\mathcal{F}_i$-measurable for all (k,ℓ) satisfying either $i \leq k$ and $i+2 \leq \ell$, or $k = \ell \in \{i, i+1\}$. Finally, let $\zeta_i \sim \mathcal{L}(M[i,i+1]|\mathcal{F}_i)$. We make the following assumptions for the matrix M:

Assumption 6.1. *(1) The random variable ζ_i satisfies the anti-concentration bound*

$$\mathbb{P}\Big[\sup_{x\in\mathbb{R}} \sup_{\beta\in\mathbb{R},(\alpha,\epsilon)\in\mathcal{R}(C)} \mathbb{P}[|\alpha(\zeta_i)^2 + \beta\zeta_i - x| < \epsilon|\mathcal{F}_i] < \frac{4C\sqrt{\epsilon}}{\sqrt{\alpha}} + n^{-2}\Big] > 1 - n^{-2} \tag{6.1}$$

for some fixed $1 \leq C < \infty$, where

$$\mathcal{R}(C) = \{(\alpha,\epsilon); \alpha > 0, \epsilon \geq (4C^2n^4)^{-1}\alpha\}.$$

(2) We have

$$\begin{aligned} \mathbb{P}[|M[n,n]| < (4Cn)^{-4}] &\leq \frac{1}{n^2}, \qquad n \text{ odd.} \\ \mathbb{P}[|M[n-1,n-1]M[n,n] - M[n-1,n]^2| < (4Cn)^{-4}] &\leq \frac{1}{n^2}, \qquad n \text{ even.} \end{aligned} \tag{6.2}$$

Remarks 6.2. *The assumption given by Inequality (6.1) is often easy to verify in practice. For example, it holds if the conditional density ρ_i of ζ_i is bounded by the constant $1 \leq C < \infty$ with high probability (see Lemma 6.4). The second part of Assumption 6.1 is often straightforward to check by hand.*

Let $|M|$ denote the determinant of the matrix M.

Lemma 6.3. *For a matrix M satisfying the hypotheses of Assumption 6.1,*

$$\mathbb{P}[|M| < (4Cn)^{-4(n+1)}] \leq \frac{3}{n}$$

for all $n \in \mathbb{N}$.

Proof. For an $n \times n$ matrix A and indices $1 \le i, j \le n$, denote by $A_{(i,j)}$ the matrix obtained by removing the i'th row and j'th column from A, and $A_{(i,j),(k,\ell)}$ the matrix obtained by removing the i'th and k'th rows and the j'th and ℓ'th columns.

Let $m = \frac{n-1}{2}$ when n is odd, and let $m = \frac{n}{2} - 1$ when n is even. Let $M^{(1)} = M$. We inductively define $\{M^{(k)}\}_{k=1}^{m}$ by setting

$$M^{(k+1)} = M^{(k)}_{(1,2),(2,1)}.$$

Observe that $M^{(k)}$ is just the $(n-2k+2)$ by $(n-2k+2)$ lower-right-hand submatrix of $M = M^{(1)}$. We define the events:

$$U_k = \{|M^{(k)}|^2 > (4Cn)^{-4(m-k+2)}\}$$
$$V_k = \Big\{ \sup_{x\in\mathbb{R}} \sup_{\beta\in\mathbb{R},(\alpha,\epsilon)\in\mathcal{R}(C)} \mathbb{P}[|\alpha(\zeta_{2k-1})^2 + \beta\zeta_{2k-1} - x| < \epsilon|\mathcal{F}_{2k-1}] < \frac{4C\sqrt{\epsilon}}{\sqrt{\alpha}} + n^{-2} \Big\}.$$

By definition $U_k \in \mathcal{F}_{2k-2}$ and $V_k \in \mathcal{F}_{2k-1}$. We now expand the determinant of $M^{(k)}$:

$$\begin{aligned} |M^{(k)}| &= -M^{(k)}[1,2]\,|M^{(k)}_{(1,2)}| + C_1 \\ &= -M^{(k)}[1,2]M^{(k)}[2,1]\,|M^{(k)}_{(1,2),(2,1)}| + M^{(k)}[1,2]C_2 + C_1 \\ &= -M^{(k)}[1,2]^2\,|M^{(k+1)}| + M^{(k)}[1,2]C_2 + C_1, \end{aligned} \tag{6.3}$$

where

$$\begin{aligned} C_1 &= \sum_{1\le j\le n,\, j\neq 2} (-1)^{j+1} M^{(k)}[1,j]\,|M^{(k)}_{(1,j)}| \\ C_2 &= \sum_{3\le j\le n} (-1)^{j} M^{(k)}[2,j]\,|M^{(k)}_{(1,2),(2,j)}|. \end{aligned}$$

By assumption, $C_1, C_2 \in \mathcal{F}_{2k-1}$. Thus from Equation 6.3 and choosing $\epsilon^2 = \frac{1}{16}C^{-4}n^{-8}|M^{(k+1)}|^2$ and $\alpha = |M^{(k+1)}|$ in Equation 6.1, we obtain

$$\mathbb{P}\Big(|M^{(k)}|^2 < \frac{1}{16}C^{-4}n^{-8}|M^{(k+1)}|^2 | V_k\Big) > 1 - \frac{2}{n^2}. \tag{6.4}$$

Using Equation (6.4) we deduce that

$$\begin{aligned} \mathbb{P}[U_k^c \cap U_{k+1} \cap V_k] &= \mathbb{E}\Big[1_{U_k^c} 1_{U_{k+1}} 1_{V_k}\Big] \\ &\le \mathbb{E}\Big[1_{U_k^c} | V_k, U_{k+1}\Big] \\ &\le \frac{4C(4Cn)^{-2(m-k+2)}}{(4Cn)^{-2(m-k+1)}} + n^{-2} \le 2n^{-2}. \end{aligned} \tag{6.5}$$

Using this inequality repeatedly, and defining $V = \cup_i V_i$, we have

$$\begin{aligned} \mathbb{P}[U_1^c] &\le \mathbb{P}[U_1^c \cap V] + \mathbb{P}[V^c] \\ &\le \mathbb{P}[U_1^c \cap V] + \frac{m}{n^2} \\ &= \mathbb{P}[U_1^c \cap U_2 \cap V] + \mathbb{P}[U_1^c \cap U_2^c \cap V] + \mathbb{P}[V^c] \\ &= \dots \end{aligned}$$

$$= \sum_{j=2}^{m} \mathbb{P}[U_1^c \cap \ldots \cap U_{j-1}^c \cap U_j \cap V] + \mathbb{P}[U_m^c] + \mathbb{P}[V^c]$$
$$\leq 2\frac{m-1}{n^2} + \frac{1}{n^2} + \frac{m}{n^2} \leq \frac{3}{n},$$

where the first few lines are repeated use of exact equalities and union bounds, and the three terms in the last line use, respectively, Inequality (6.5), Inequality (6.2) and the assumption in Inequality (6.1) that $\mathbb{P}[V^c] \leq n^{-2}$. This completes the proof. □

We give a simple sufficient condition for the assumption given by Inequality (6.1) to hold:

Lemma 6.4 (Non-Concentration)**.** *Define* $f : \mathbb{R} \mapsto \mathbb{R}$ *by*

$$f(x) = \alpha x^2 + \beta x + \gamma$$

and let X *be a random variable with density* ρ *satisfying* $\sup_x \rho(x) < C < \infty$*. Then for any* $x \in \mathbb{R}$ *and* $\epsilon > 0$*,*

$$\mathbb{P}[|f(X) - x| < \epsilon] \leq \frac{4C\sqrt{\epsilon}}{\sqrt{|\alpha|}}.$$

Proof. Assume without loss of generality that $\alpha > 0$, and let η be the density of $f(X)$. Fix $r \in \mathbb{R}$ and define the quantity

$$\Gamma = \frac{r}{\alpha} + \frac{\beta^2}{4\alpha^2} - \frac{\gamma}{\alpha}.$$

We have

$$\mathbb{P}[f(X) \leq r] = \mathbb{P}[(X - \frac{\beta}{2\alpha})^2 \leq \Gamma]$$
$$= \mathbb{P}[X \leq \frac{\beta}{2\alpha} + \sqrt{\Gamma}] - \mathbb{P}[X \leq \frac{\beta}{2\alpha} - \sqrt{\Gamma}].$$

Thus

$$\eta(r) = \frac{d}{dr}\mathbb{P}[f(X) \leq r]$$
$$= \rho(\frac{\beta}{2\alpha} + \sqrt{\Gamma})\frac{1}{2\alpha\sqrt{\Gamma}} + \rho(\frac{\beta}{2\alpha} - \sqrt{\Gamma})\frac{1}{2\alpha\sqrt{\Gamma}}$$
$$\leq \frac{C}{\alpha}\frac{1}{\sqrt{\Gamma}}.$$

Thus, we have

$$\mathbb{P}[|f(X) - x| \leq \epsilon] = \int_{x-\epsilon}^{x+\epsilon} \eta(r)dr$$
$$\leq \frac{C}{\alpha}\int_{x-\epsilon}^{x+\epsilon} \frac{1}{\sqrt{\Gamma}}$$
$$\leq \frac{C}{\alpha}\int_{-\epsilon}^{\epsilon} \frac{\sqrt{\alpha}}{\sqrt{|r|}}dr$$
$$= \frac{4C\sqrt{\epsilon}}{\sqrt{\alpha}}.$$

This completes the proof of the bound. □

We next show that the assumption given by Inequality (6.1) remains true under small perturbations. For two probability measures ν_1, ν_2 on $\mathbb{R}^d$, define the *Wasserstein distance*

$$W_2(\nu_1, \nu_2)^2 = \inf_{(X,Y)\in\mathcal{C}, X\sim\nu_1, Y\sim\nu_2} \mathbb{E}[\|X - Y\|_2^2],$$

where $\mathcal{C}$ is the set of all couplings on $\mathbb{R}^d \times \mathbb{R}^d$ with marginal distributions ν_1 and ν_2.

Let M' be another $n \times n$ symmetric random matrix, and let $\mathcal{F}_i'$, ζ_i', etc be defined analogously to $\mathcal{F}_i$, ζ_i. For fixed $1 \le i \le n$, let $\mathcal{G}_i = \mathcal{G}_i(\{m_{k,\ell}\})$ (respectively $\mathcal{G}_i'(\{m_{k,\ell}\})$) be the event that $M[k,\ell] = m_{k,\ell}$ (respectively $M'[k,\ell] = m_{k,\ell}$) for all $k, \ell \ge i+2$.

Lemma 6.5. *Let the density ρ_i of ζ_i satisfies*

$$\mathbb{P}[\sup_x \rho_i(x) > C] \le \frac{1}{2n^5} \tag{6.6}$$

for some fixed $1 \le C < \infty$. Let M' satisfy (6.2) of Assumption 6.1. Assume that

$$W_2(\mathcal{L}(M[i,(i+1):n] \,|\, \mathcal{G}_i), \mathcal{L}(M'[i,(i+1):n] \,|\, \mathcal{G}_i'))^2 \le \delta < \frac{1}{8n^5}(4Cn)^{-4}. \tag{6.7}$$

Then there exists a universal constant N_0 so that the determinant $|M'|$ of M' satisfies:

$$\mathbb{P}[|M'| < (4Cn^3)^{-4(n+1)}] \le \frac{3}{n}$$

for all $n > N_0$.

Proof. We begin with a generic bound. Let $X, Y \in \mathbb{R}$ be two random variables, and let $f : \mathbb{R} \mapsto \mathbb{R}$ by $f(x) = \alpha x^2 + \beta x$ for some $\alpha > 0$. Fix $\epsilon > 0$ and $x \in \mathbb{R}$. Then

$$\begin{aligned}
\sup_{x\in\mathbb{R}} \mathbb{P}[|f(X) - x| < \epsilon] &= \sup_{x\in\mathbb{R}} \mathbb{P}[|\alpha X^2 + \beta X - x| \le \epsilon] \\
&\le \sup_{x\in\mathbb{R}} \mathbb{P}\Big[X \in \Big(x - \sqrt{\frac{\epsilon}{\alpha}}, x + \sqrt{\frac{\epsilon}{\alpha}}\Big)\Big] \\
&\le \sup_{x\in\mathbb{R}} \mathbb{P}[Y \in [x - \sqrt{\frac{\epsilon}{4\alpha}}, x + \sqrt{\frac{\epsilon}{4\alpha}}]] + \mathbb{P}[|X - Y| > \sqrt{\frac{\epsilon}{4\alpha}}] \\
&\le \sup_{x\in\mathbb{R}} \mathbb{P}\Big[Y \in \Big(x - \sqrt{\frac{\epsilon}{4\alpha}}, x + \sqrt{\frac{\epsilon}{4\alpha}}\Big)\Big] + W_2(\mathcal{L}(X), \mathcal{L}(Y))^2 \frac{4\alpha}{\epsilon}
\end{aligned} \tag{6.8}$$

By Equation (6.6) and Lemma 6.4, Assumption 6.1 is satisfied the by the matrix M. For ϵ, f as above, Equation (6.8) then implies

$$\begin{aligned}
&\sup_{x\in\mathbb{R}} \mathbb{P}[f(M'[i,i+1]) \in [x-\epsilon, x+\epsilon] \,|\, \mathcal{G}_i'] \\
&\quad\le \sup_{x\in\mathbb{R}} \mathbb{P}[f(M[i,i+1]) \in [x-\epsilon, x+\epsilon] \,|\, \mathcal{G}_i] \\
&\qquad\qquad + \frac{4\alpha}{\epsilon} W_2\Big(\mathcal{L}(M[i,(i+1):n] \,|\, \mathcal{G}_i), \mathcal{L}(M'[i,(i+1):n] \,|\, \mathcal{G}_i')\Big)^2 \\
&\quad\le \frac{1}{2n^5} + \frac{4C\sqrt{\epsilon}}{\alpha} + \frac{4\alpha}{\epsilon} W_2\Big(\mathcal{L}(M[i,(i+1):n] \,|\, \mathcal{G}_i), \mathcal{L}(M'[i,(i+1):n] \,|\, \mathcal{G}_i')\Big)^2 \\
&\quad\le \frac{1}{2n^5} + \frac{4C\sqrt{\epsilon}}{\alpha} + \frac{4\alpha\delta}{\epsilon},
\end{aligned}$$

where the first term in the second-last line comes from Inequality (6.6). and the second term comes from an application of Lemma 6.4, and where the last line comes from Inequality (6.7). This bound implies

$$\mathbb{E}[\sup_{x\in\mathbb{R}} \mathbb{P}[f(M'[i,i+1]) \in [x-\epsilon, x+\epsilon] \,|\, \mathcal{F}'_i] \,|\, \mathcal{G}'_i] \leq \frac{1}{2n^5} + \frac{4C\sqrt{\epsilon}}{\alpha} + \frac{4\alpha\delta}{\epsilon} \leq \frac{1}{n^5} + \frac{4C\sqrt{\epsilon}}{\alpha}.$$

By Markov's inequality, this implies

$$\mathbb{P}[\sup_{x\in\mathbb{R}} \mathbb{P}[f(M'[i,i+1]) \in [x-\epsilon, x+\epsilon] \,|\, \mathcal{F}'_i] > n^2(n^{-5} + \frac{4C\sqrt{\epsilon}}{\alpha})] \leq n^{-2}.$$

Thus M' satisfies Assumption 6.1 with constant $\tilde{C} = n^2 C$. Applying Lemma 6.3 now completes the proof. □

6.2. **Bounding the Smallest Singular Value of D_∞.** In this section, we define a specific random matrix D_∞, and show that it satisfies the requirements of Lemma 6.3. Define the n-sphere

$$\mathrm{S}^{(n-1)} = \{x \in \mathbb{R}^n \,:\, \sum_{i=1}^{n} x[i]^2 = 1\}.$$

Denote the Euclidean inner product by $\langle \cdot, \cdot \rangle$. We begin with the technical lemma:

Lemma 6.6 (Conditional Densities on Spheres)**.** *Fix $0 \leq k \leq n-2$. Let $X, v_1, \ldots, v_{n-1} \sim \mathrm{Unif}(\mathrm{S}^{(n-1)})$ be i.i.d. For $0 \leq k \leq n-2$, let $\mathcal{G}_k$ be the σ-algebra generated by $\langle X, v_1\rangle, \ldots, \langle X, v_k\rangle$ and $v_1, \ldots, v_{n-1}$. Let*

$$Z \sim \mathcal{L}(\langle X, v_{n-1}\rangle \,|\, \mathcal{G}_k),$$

and let ρ be the density of Z. We have

$$\mathbb{P}[\sup_z \rho(z) > n^{20}] \leq n^{-2}$$

for all $n > N_0$ sufficiently large, uniformly in $0 \leq k \leq n-2$.

Proof. We first prove the claim for $k = n-2$, the maximum value.Let $\mathcal{F} = \mathcal{G}_{n-2}$. We can write down the dependence of $\sup_z \rho(z)$ on $\mathcal{F}$ explicitly. Let $H = \mathrm{span}(v_1, \ldots, v_{n-2})$ be the hyperplane spanned by $v_1, \ldots, v_{n-2}$. Let $\mathcal{P}_H : \mathbb{R}^n \mapsto H$ be the operator associated with orthogonal projection onto H, and define

$$v^0 = \mathcal{P}_H(v_{n-1})$$
$$v^+ = v_{n-1} - \mathcal{P}_H(v_{n-1}).$$

We note that $\langle X, v^0\rangle$ and v^+ are both $\mathcal{F}$-measurable. Let $X_H = \|\mathcal{P}_H X\|_2^2$. The random variable X_H is also $\mathcal{F}$-measurable.

Let $Z^+ = \langle X, v^+\rangle$, so that

$$\mathcal{L}(Z) = \mathcal{L}(Z^+ + \langle X, v^0\rangle | \mathcal{G}_k). \tag{6.9}$$

Let $Y = \|v^+\|_2 \sqrt{1 - X_H}\|S$, where $S \sim \mathrm{Unif}(\mathrm{S}^{(1)})$. Then $Z^+ \stackrel{D}{=} Y[1]$. Thus, the density ρ^+ of Z^+ satisfies

$$\begin{aligned}\rho^+(z) &= \frac{2}{\pi}\frac{\sqrt{\|v^+\|_2^2(1-X_H) - z^2}}{\|v^+\|_2^2(1-X_H)} \\ &\leq \frac{2}{\pi}\frac{1}{\sqrt{\|v^+\|_2^2(1-X_H)}}.\end{aligned}$$

By Equation 6.9 and the fact that $\langle X, v^0\rangle \in \mathcal{F}$, it follows that

$$\sup_z \rho(z) = \sup_z \rho^+(z) \leq \frac{2}{\pi}\frac{1}{\sqrt{\|v^+\|_2^2(1-X_H)}}. \tag{6.10}$$

Thus, to complete the proof, it suffices to bound $\frac{1}{\sqrt{\|v^+\|_2^2(1-X_H)}}$ with high probability.

First, we bound $1 - X_H$. Since $X_H = \|\mathcal{P}_H\|_2^2$ where $X \sim \mathrm{Unif}(\mathrm{S}^{(n-1)})$ and H is an independently and randomly chosen hyperplane of dimension $n-2$, we can assume without loss of generality that $H = \{x \in \mathbb{R}^n : x[n-1] = x[n] = 0\}$. Thus,

$$\begin{aligned}\mathbb{P}[1 - X_H \leq n^{-20}] &= \mathbb{P}[1 - X[1]^2 - \ldots - X[n-2]^2 \leq n^{-20}] \\ &= \mathbb{P}[X[n-1]^2 + X[n]^2 \leq n^{-20}] \\ &= O(n^{-15}).\end{aligned} \tag{6.11}$$

By exactly the same reasoning, we deduce that

$$\begin{aligned}\mathbb{P}[\|v^+\|_2^2 \leq n^{-5}] &= \mathbb{P}[X[n]^2 \leq n^{-20}] \\ &= O(n^{-15}).\end{aligned} \tag{6.12}$$

Combining Inequalities (6.11) and (6.12) with Inequality (6.10), we conclude that

$$\mathbb{P}[\sup_z \rho(z) > n^{-20}] = O(n^{-15}).$$

This completes the proof of the lemma. □

Let $P_1, \ldots, P_N$ be i.i.d. draws from the Haar measure on $\mathrm{SO}(N)$. Define a symmetric $N \times N$ matrix D_∞ by

$$\begin{aligned}D_\infty[i,i] &= 1, \\ D_\infty[i,j] &= -\mathrm{Tr}[a_i(\prod_{\ell=i+1}^{j} P_\ell)a_j(\prod_{\ell=i+1}^{j} P_\ell)^{-1}],, \qquad i < j \\ D_\infty[i,j] &= D_\infty[j,i], \qquad i > j.\end{aligned} \tag{6.13}$$

For $i < j$, let $P_{i,j} = \prod_{\ell=i+1}^{j} P_\ell$. We note that, for $1 \leq i < j \leq n$,

$$D_\infty[i,j] = -\mathrm{Tr}[a_i P_{i,j} a_j (P_{i,j})^{-1}].$$

Remark 6.7. The matrix D_∞ has two useful properties. First, the appearance of the Haar measure in the definition of D_∞ means that it is easier to make exact calculations involving the entries of D_∞ than those of D. Second, D_∞ is 'close' to D, and so bounds on D_∞ can easily be transferred to bounds on D. More precisely, Theorem 1 of [Oli09] implies that the

entries of D converge to those of D_∞ as $\mathcal{Q}$ goes to infinity. The $\Theta(n^2\log(n))$ scaling of the lower bound $s_{i+1} - s_i \geq \mathcal{Q}n^2\log(n)$ in Definition 2.6 was chosen so that we could guarantee that D and D_∞ must be close in distribution.

For $1 \leq i < N$, we define $\mathcal{F}_i$ to be the σ-algebra generated by the matrices $\{P_j\}$ and the inner products $\{-\mathrm{Tr}[a_i P_{i,j} a_j (P_{i,j})^{-1}]\}$ for $j > i+1$.

Lemma 6.8. *The matrix D_∞ and σ-algebras $\{\mathcal{F}_i\}_{1\leq i<n}$ given above satisfy Assumption 6.1 with*

$$C = N^{20}$$

for all $N > N_0$ sufficiently large. Thus for all $N > N_0$ sufficiently large,

$$\mathbb{P}[|D_\infty| < (4N^{21})^{-4(N+1)}] \leq \frac{3}{N}.$$

Proof. We make some initial observations. Let

$$\mathfrak{S}_n = \{a \in \mathfrak{so}(n) \,:\, \|a\|_{\mathrm{HS}} = 1\},$$

and define the bijection $\mathcal{M} : \mathfrak{S}_n \mapsto \mathrm{S}^{(N-1)}$ by

$$\mathcal{M}(a) = (\langle a, a_1\rangle_{\mathrm{HS}}, \ldots, \langle a, a_N\rangle_{\mathrm{HS}}).$$

We note that, if $P \sim \mathrm{Unif}(\mathrm{SO}(n))$ and $a \in \mathfrak{S}_n$, then

$$\mathcal{M}(PaP^{-1}) \sim \mathrm{Unif}(\mathrm{S}^{(N-1)}). \tag{6.14}$$

We first prove Inequality (6.2). When N is odd, this is trivial. When N is even, we let $X \sim \mathrm{Unif}(\mathrm{S}^{(n-1)})$. By Equation (6.14), we have that $D_\infty[N-1, N] \overset{D}{=} X[1]$, so

$$\begin{aligned}\mathbb{P}[|D_\infty[N,N]D_\infty[N-1,N-1] - D_\infty[N-1,N]^2| < (4N^6)^{-4}] &= \mathbb{P}[|1 - X[1]^2| < (4N^6)^{-4}] \\ &= O(N^{-3}).\end{aligned}$$

We next prove Inequality (6.1). Observe that:

(1) P_{i+1} is independent of $\{P_j\}_{j>i+1}$.

(2) By Equality (6.14), the vectors $\{\mathcal{M}((P_{i+2}\ldots P_j)a_j(P_{i+2}\ldots P_j)^{-1})\}_{j>i+1}$ are i.i.d. choices from the sphere $S^{(N-1)}$, and the entries $\{D_\infty[i,j]\}_{j>i+1}$ are inner products of these vectors with $\mathcal{M}(a_i)$:

$$\begin{aligned}D_\infty[i, j+1] &= -\mathrm{Tr}[a_i(\prod_{\ell=i+1}^{j} P_\ell)a_j(\prod_{\ell=i+1}^{j} P_\ell)^{-1}] \\ &= -\langle \mathcal{M}(a_i), \mathcal{M}((P_{i+2}\ldots P_j)a_j(P_{i+2}\ldots P_j)^{-1})^\dagger\rangle.\end{aligned}$$

(3) Combining the previous two observations, the distribution of

$$Z \sim \mathcal{L}(-\langle \mathcal{M}(a_i), \mathcal{M}(P_{i+1}a_{i+1}P_{i+1}^{-1})\rangle | \mathcal{F}_i)$$

satisfies the requirements of Lemma 6.6.

Thus, by Lemma 6.4, the matrix D_∞ satisfies Inequality (6.1) of Assumption 6.1 with constant $C = N^{20}$. The conclusion follows immediately from Lemma 6.3. This completes the proof. □

We apply our results to obtain a bound on the smallest singular value of D_∞:

Lemma 6.9 (Smallest Singular Values of D_∞)**.** *Let D_∞ be the matrix defined as in (6.13), and let $\sigma_1(D_\infty) \le \sigma_2(D_\infty) \le \ldots \le \sigma_N(D_\infty)$ be its singular values. Then*

$$\mathbb{P}[\sigma_1(D_\infty) \le N^{-N}(4N^{21})^{-4(N+1)}] = o(1).$$

Proof. Using the trivial bound $\sigma_i(D_\infty) \le N \max_{k,\ell} |D_\infty[k,\ell]| \le N$, we have

$$\begin{aligned}\sigma_1(D_\infty) &= |D_\infty| \prod_{i=2}^{n} \sigma_i(D_\infty)^{-1} \\ &\ge |D_\infty|\, N^{-N}.\end{aligned}$$

Thus, for any $0 < r < \infty$.

$$\mathbb{P}[\sigma_1(D_\infty) \le r] \le \mathbb{P}[|D_\infty| \le rN^N].$$

Choosing $r = N^{-N}(4N^{21})^{-4(N+1)}$ and applying Lemmas 6.3 and 6.8, we have

$$\begin{aligned}\mathbb{P}[\sigma_1(D_\infty) \le N^{-N}(4N^{21})^{-4(N+1)}] &\le \mathbb{P}[|D_\infty| \le N^{-N}(4N^{21})^{-4(N+1)}N^N] \\ &= \mathbb{P}[|D_\infty| \le (4N^{21})^{-4(N+1)}] \\ &\le \frac{3}{N}.\end{aligned}$$

This completes the proof. □

6.3. **Application to Kac's Walk.** We show that the sequence $\{\phi_n^{-1}\}_{n\ge1}$ defined in Equation (5.1) does not grow too quickly, and thus complete our proof of Theorem 1.1. We do this by comparing the matrix D of interest, defined in Equation (3.6), with the matrix D_∞ studied in Section 6.2.

Lemma 6.10. *Let $\mathcal{Q} = 100$. Then the sequence $\{\phi_n^{-1}\}_{n\ge1}$ defined in Equation (5.1) and associated with the coupling defined in Section 2.1 satisfies*

$$\phi_n^{-1} \le (4N^{24})^{5N}$$

for all $n > N_0$ sufficiently large.

Remark 6.11. We believe that Lemma 6.10 may hold with $\phi_n = O(n^k)$ for some $k < \infty$.

Proof. We begin by relating the symmetric matrices D, D_∞. Recall the law of the off-diagonal entries of D given in Equation (3.6):

$$D[i,j] = -\mathrm{Tr}[a_i(\prod_{\ell=i+1}^{j} M_\ell)a_j(\prod_{\ell=i+1}^{j} M_\ell)^{-1}], \qquad i<j, \tag{6.15}$$

where

$$M_\ell = \prod_{t=s_\ell+1}^{s_{\ell+1}} R(i_s, \theta_s).$$

We will relate D to D_∞. Recall that the off-diagonal entries of D_∞ are written

$$D_\infty[i,j] = -\mathrm{Tr}[a_i(\prod_{\ell=i+1}^{j} M'_\ell)a_j(\prod_{\ell=i+1}^{j} M'_\ell)^{-1}], \qquad i<j,$$

where $\{M_\ell\}$ are an i.i.d. sequence from the Haar measure on $\mathrm{SO}(n)$. For fixed $1 \le i < N$, let $\mathcal{G}_i = \mathcal{G}_i(\{m_{k,\ell}\})$ (respectively $\mathcal{G}'_i(\{m_{k,\ell}\})$) be the event that $D[k,\ell] = m_{k,\ell}$ (respectively $D_\infty[k,\ell] = m_{k,\ell}$) for all $k, \ell \ge i+2$. By the main result of [Oli09] and Lemma 3.1, we have

$$W_2(\mathcal{L}(D[i,(i+1):n] \,|\, \mathcal{G}_i), \mathcal{L}(D_\infty[i,(i+1):n] \,|\, \mathcal{G}'_i))^2 \le 2n^{-\mathcal{Q}+3} \tag{6.16}$$

for all $\mathcal{Q} \ge 5$ and all $n > N_0$ sufficiently large.

With this initial calculation complete, we can now prove Lemma 6.10. Fix notation as in Lemma 6.16. We will apply Lemma 6.5, with D_∞ playing the role of M and D playing the role of M'. By Lemma 6.8, D_∞ satisfies the conditions of Lemma 6.5 with constant $C = N^{20}$. By Inequality 6.16, D_∞, D satisfy Condition (6.7) of Lemma 6.5 for all $\mathcal{Q} > 94$. Thus, for fixed $\mathcal{Q} > 94$,

$$\mathbb{P}[|D| > (4N^{23})^{5N}] \le \frac{3}{n}$$

for all $n > N_0$ sufficiently large. By a calculation identical to that in Lemma 6.9, we conclude that

$$\mathbb{P}[\sigma_1(D) > (4N^{24})^{5N}] \le \frac{3}{n},$$

completing the proof. □

7. Proof of Theorem 1.1

Inequality (1.4), the lower bound in Theorem 1.1, is exactly the main conclusion of Theorem 1.

To prove Inequality (5.5), the upper bound in Theorem 1.1, we recall from Lemma 6.10 that the sequence $\{\phi_n^{-1}\}_{n\ge1}$ defined in Equation (5.1) satisfies

$$\phi_n^{-1} \le (4N^{24})^{5N} = o(n^{n^{2.2}})$$

whenever $\mathcal{Q} > 100$.

Thus, in Theorem 2, we may take $\mathcal{Q} = 101$, $R = 808n^4 \log(n)$ and $\log(\phi_n) \le 120\, n^2\, (\log(4) + \log(n))$. Thus, by Theorem 2, we have

$$\limsup_{n\to\infty} \sup_{X_0 = x \in \mathrm{SO}(n)} \|\mathcal{L}(X_t) - \mu\|_{\mathrm{TV}} = 0$$

for any sequence $t = t(n)$ satisfying

$$t(n) \ge 10^7\, n^4 \log(n) \ge 808n^4 \log(n) + 5n^2 \log(n) + 900n^2 \times (120\, n^2\, (\log(4) + \log(n))).$$

This completes the proof of Inequality (5.5).

8. Discussion

Our work leaves open the question as to whether the mixing time of Kac's walk is indeed $\Theta(n^2 \log(n))$ as conjectured, and whether it exhibits the *cutoff phenomenon*. We cannot obtain such a bound by more careful analysis of the terms in Theorem 2. However, we believe that it is possible to obtain the desired bound by a more careful analysis of the terms in Theorem 3.

The main difficulty in applying our method is obtaining a bound on ϕ_n, which measures the smallest singular value $\sigma_1(D)$ of the matrix D. In this paper, we were only able to bound $\sigma_1(D)$ by comparing D to a simpler limiting matrix D_∞ for which exact calculations

were available. To improve this bound further, we believe that it is necessary to analyze $\sigma_1(D)$ directly. To obtain an $O(n^2 \log(n))$ bound on the mixing time of Kac's walk using our method, it would be enough to obtain any polynomial bound $\phi_n = O(n^k)$ for some $0 < k < \infty$. The main obstacles to proving such a bound are:

(1) Our weak random matrix bound in Lemma 6.3. Our argument for Lemma 6.3, like that in [FG13], only takes advantage of the randomness of entries within distance 1 of the diagonal of an n by n random matrix M. Unfortunately, no argument that only analyzes these entries can give any bound that is stronger than $|M|^{-2} = 2^{O(n)}$. Since the matrix D of interest has 'many more' than $3n$ 'pieces' of randomness, we can hope to take advantage of them and obtain a stronger bound, as in [FV15].

(2) We bound the determinant $|D|$ of the random matrix D and use this to obtain a bound on the smallest singular value $\sigma_1(D)$; we give up a factor of N^{N-1} in the process (see Lemma 6.9). To avoid this loss, we must either obtain stronger bounds on the joint distribution of the remaining singular values $\sigma_2(D) \leq \sigma_3(D) \leq \ldots \leq \sigma_N(D) \gtrsim \log(n)$, or to bound $\sigma_1(D)$ directly as in [FV15].

(3) Our bound on $\sigma_1(D)$ is obtained through a comparison of D and the related matrix D_∞ (see Theorem 6.10). While D, D_∞ are nearby under the coupling studied in Theorem 2, they are very far under the coupling studied in Theorem 3. Thus, to get a better mixing bound, the matrix D should be studied directly.

It seems possible to extend our arguments to resolve any two of these three obstacles together; however, we see no route to resolving all three simultaneously.

Our approach can be applied to the analysis of other Gibbs samplers on constrained or non-convex state spaces. As mentioned in Remark 2.8, our non-Markovian coupling can be defined for generic Gibbs samplers. The key ingredients in an analysis of the entire coupling are the analysis of an underlying 'scaffolding' coupling $\{\hat{X}_t, \hat{Y}_t\}_{t \geq 0}$, the analysis of the smallest singular value of the Jacobian of the 'perturbation map' $f_{\mathcal{A},\epsilon_n}$, and soft bounds on the smoothness of $f_{\mathcal{A},\epsilon_n}$.

Acknowledgements

NSP is partially supported by an ONR grant and AMS is partially supported by an NSERC grant. We thank Tristan Collins, Persi Diaconis, John Jiang, Federico Poloni, Terence Tao and Roman Vershynin for helpful conversations and suggestions.

References

[AC06] Nir Ailon and Bernard Chazelle. Approximate nearest neighbors and the fast Johnson-Lindenstrauss transform. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 557–563. ACM, 2006.

[AKS82] Miklos Ajtai, Janos Komlos, and Endre Szemerédi. Largest random component of a k-cube. *Combinatorica*, 2(1):1–7, 1982.

[AM] Hossein Abbaspour and Martin A Moskowitz. *Basic Lie Theory*. World Scientific.

[Bro] Daniel Brown. How I wasted too long finding a concentration inequality for sums of geometric variables. https://cs.uwaterloo.ca/~browndg/negbin.pdf. Accessed: 2016-04-29.

[Cap08] Pietro Caputo. On the spectral gap of the Kac walk and other binary collision processes. *arXiv preprint arXiv:0807.3415*, 2008.

[CCL03] Eric A Carlen, Maria C Carvalho, and Michael Loss. Determination of the spectral gap for Kac's master equation and related stochastic evolution. *Acta mathematica*, 191(1):1–54, 2003.

[CCR+08] Eric A Carlen, Maria C Carvalho, J Le Roux, Michael Loss, and Cédric Villani. Entropy and chaos in the Kac model. *arXiv preprint arXiv:0808.3192*, 2008.

[CF14] Roberto Cortez and Joaquin Fontbona. Quantitative propagation of chaos for generalized Kac particle systems. *arXiv preprint arXiv:1406.2115*, 2014.

[CG92] George Casella and Edward I George. Explaining the Gibbs sampler. *The American Statistician*, 46(3):167–174, 1992.

[CV06] Mark Conger and Divakar Viswanath. Shuffling cards for blackjack, bridge, and other card games. *arXiv preprint math/0606031*, 2006.

[CV14] Artur Czumaj and Berthold Vöcking. Thorp shuffling, butterflies, and non-Markovian couplings. In *Automata, Languages, and Programming*, pages 344–355. Springer, 2014.

[Dia88] Persi Diaconis. Group representations in probability and statistics. *Lecture Notes-Monograph Series*, pages i–192, 1988.

[DKSC08] Persi Diaconis, Kshitij Khare, and Laurent Saloff-Coste. Gibbs sampling, exponential families and orthogonal polynomials. *Statistical Science*, 23(2):151–178, 2008.

[DSC00] Persi Diaconis and Laurent Saloff-Coste. Bounds for Kac's master equation. *Communications in Mathematical Physics*, 209(3):729–755, 2000.

[ER61] Paul Erdös and Alfréd Rényi. On a classical problem of probability theory. *Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei*, 6:215–220, 1961.

[FG13] Omer Friedland and Ohad Giladi. A simple observation on random matrices with continuous diagonal entries. *Electronic Communications in Probability*, 18(53):1–7, 2013.

[Fur02] Alex Furman. Random walks on groups and random transformations. *Handbook of dynamical systems*, 1:931–1014, 2002.

[FV15] Brendan Farrell and Roman Vershynin. Smoothed analysis of symmetric random matrices with continuous distributions. *Proceedings of the American Mathematical Society*, 2015.

[GP74] V. Guilleman and A. Pollack. *Differential Topology*. Prentice Hall, 1974.

[GS01] Charles Goldman, Herbert an Poole and John Safko. *Classical Mechanics (3rd Edition)*. Pearson, 2001.

[Has70] W Keith Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

[HJ16] Bob Hough and Yunjiang Jiang. Cut-off phenomenon in the uniform plane Kac walk. *Annals of Probability (Forthcoming)*, 2016.

[HM14] Maxime Hauray and Stéphane Mischler. On Kac's chaos and related problems. *Journal of Functional Analysis*, 266(10):6055–6157, 2014.

[Hog14] Leslie Hogben. *Handbook of linear algebra*. 2nd edition, 2014.

[HV03] Thomas P Hayes and Eric Vigoda. A non-Markovian coupling for randomly sampling colorings. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 618–627. IEEE, 2003.

[Jam55] AT James. A generating function for averages over the orthogonal group. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 229(1178):367–375, 1955.

[Jan01] Élise Janvresse. Spectral gap for Kac's model of Boltzmann equation. *The Annals of Probability*, 29(1):288–304, 2001.

[Jan03] Élise Janvresse. Bounds on semigroups of random rotations on SO(n). *Theory of Probability & Its Applications*, 47(3):526–532, 2003.

[JH01] Galin L Jones and James P Hobert. Honest exploration of intractable probability distributions via Markov chain Monte Carlo. *Statistical Science*, pages 312–334, 2001.

[Jia12a] Yunjiang Jiang. Total variation bound for Kac's random walk. *The Annals of Applied Probability*, 22(4):1712–1727, 2012.

[Jia12b] Yunjiang Jiang. Total variation mixing time of Kac's master walk. *Preprint*, 2012.

[Kac54] Mark Kac. Foundations of kinetic theory. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, number 3, pages 171–197, 1954.

[Lov99] László Lovász. Hit-and-run mixes fast. *Mathematical Programming*, 86(3):443–461, 1999.

[LV03] László Lovász and Santosh Vempala. Hit-and-run is fast and fun. 2003.

[Mas03] David K Maslen. The eigenvalues of Kac's master equation. *Mathematische Zeitschrift*, 243(2):291–331, 2003.

[MM13] Stéphane Mischler and Clément Mouhot. Kac's program in kinetic theory. *Inventiones mathematicae*, 193(1):1–147, 2013.

[Oli09] Roberto Imbuzeiro Oliveira. On the convergence to equilibrium of Kac's random walk on matrices. *The Annals of Applied Probability*, pages 1200–1231, 2009.

[Por96] Ursula Porod. The cut-off phenomenon for random reflections. *The Annals of Probability*, 24(1):74–96, 1996.

[PS07] Igor Pak and Sergiy Sidenko. Convergence of Kac's random walk. *Preprint available from http://www-math.mit.edu/˜pak/research.html*, 2007.

[PS15] Natesh S Pillai and Aaron Smith. Kac's walk on n-sphere mixes in $n \log n$ steps. *arXiv preprint arXiv:1507.08554*, 2015.

[Ray02] Jean Saint Raymond. Local inversion for differentiable functions and the Darboux property. *Mathematika*, 49(1-2):141–158, 2002.

[Ros94] Jeffrey S Rosenthal. Random rotations: characters and random walks on SO(n). *The Annals of Probability*, pages 398–423, 1994.

[SC04] Laurent Saloff-Coste. Random walks on finite groups. In *Probability on discrete structures*, pages 263–346. Springer, 2004.

[Smi14] Aaron Smith. A Gibbs sampler on the n-simplex. *The Annals of Applied Probability*, 24(1):114–130, 2014.

Appendix A. Proofs of Technical Bounds

We prove the bounds in Section 3.

A.1. **Matrix Estimates.**

Proof of Lemma 3.1. We note the 'telescoping sum' identity

$$\prod_{i=1}^{k} Q_i - \prod_{i=1}^{k} P_i = \sum_{i=1}^{k} (\prod_{\ell=1}^{i-1} Q_\ell)(Q_i - P_i)(\prod_{\ell=i+1}^{k} P_\ell).$$

The result then follows immediately from application of the triangle inequality and the inequality $\|ABC\|_{\mathrm{HS}} \leq \|A\|_{\mathrm{Op}}\|B\|_{\mathrm{HS}}\|C\|_{\mathrm{Op}}$ for any $A, B, C \in M(n)$. □

Proof of Lemma 3.2. By Fact 4 in chapter 15 of [Hog14],

$$\sum_{i=1}^{n} (\sigma_i(M_1) - \sigma_i(M_2))^2 \leq \sum_{i=1}^{n} \sigma_i(M_1 - M_2)^2.$$

By assumption (3.2), this implies

$$\sum_{i=1}^{n} (\sigma_i(M_1) - \sigma_i(M_2))^2 \leq N\delta^2 \sigma_1(M_1)^2.$$

In particular,

$$\max_{1 \leq i \leq N} |\sigma_i(M_1) - \sigma_i(M_2)| \leq \sqrt{N}\delta\sigma_1(M_1).$$

For a symmetric matrix M, let $\lambda_1(M), \ldots, \lambda_N(M)$ denote the eigenvalues, ordered so that $|\lambda_i(M) = \sigma_i(M)$. We have

$$\begin{aligned}
|\frac{\det(M_2)}{\det(M_1)} - 1| &= |\frac{\prod_{i=1}^{N} \lambda_i(M_2)}{\prod_{i=1}^{N} \lambda_i(M_1)} - 1| \\
&= |\prod_{i=1}^{N} \frac{\lambda_i(M_1) + (\lambda_i(M_2) - \lambda_i(M_1))}{\lambda_i(M_1)} - 1| \\
&= \prod_{i=1}^{N} |\frac{(\lambda_i(M_2) - \lambda_i(M_1))}{\lambda_i(M_1)}| \\
&\leq \prod_{i=1}^{N} |\frac{\sqrt{N}\delta\sigma_1(M_1)}{\sigma_1(M_1)} \\
&= N^{\frac{N}{2}}\delta^N
\end{aligned}$$

and the proof is finished. □

Proof of Lemma 3.3. Let $x = (x_1, \ldots, x_N), y = (y_1, \ldots, y_N)$. By Equation (3.4) it follows that

$$\begin{aligned}
&\|\mathrm{d}f_x(h) - \mathrm{d}f_y(h)\|_{\mathrm{HS}} \\
&\quad = \|\sum_{j=1}^{N} \prod_{k=1}^{N} \mathrm{R}_k\, e^{(\theta_k + x_k)a_k}(\mathrm{Id} + (h_k a_k - \mathrm{Id})1_{k=j}) - \sum_{j=1}^{N} \prod_{k=1}^{N} \tilde{\mathrm{R}}_k\, e^{(\tilde{\theta}_k + y_k)a_k}(\mathrm{Id} + (h_k a_k - \mathrm{Id})1_{k=j})\|_{\mathrm{HS}}
\end{aligned}$$

$$
\begin{aligned}
&\leq \sum_{j=1}^{N} \| \prod_{k=1}^{N} \mathrm{R}_k \, e^{(\theta_k + x_k)a_k} (\mathrm{Id} + (h_k a_k - \mathrm{Id}) 1_{k=j}) - \prod_{k=1}^{N} \tilde{\mathrm{R}}_k \, e^{(\tilde{\theta}_k + y_k)a_k} (\mathrm{Id} + (h_k a_k - \mathrm{Id}) 1_{k=j}) \|_{\mathrm{HS}} \\
&\leq \sum_{j,k=1}^{N} \| \mathrm{R}_k \, e^{(\theta_k + x_k)a_k} (\mathrm{Id} + (h_k a_k - \mathrm{Id}) 1_{k=j}) - \tilde{\mathrm{R}}_k \, e^{(\tilde{\theta}_k + y_k)a_k} (\mathrm{Id} + (h_k a_k - \mathrm{Id}) 1_{k=j}) \|_{\mathrm{HS}} \\
&\leq 2N^2 \max(\|R_k - \tilde{R}_k\|_{\mathrm{HS}}, \| e^{(\theta_k + x_k)a_k} - e^{(\tilde{\theta}_k + y_k)a_k} \|_{\mathrm{HS}}) \\
&\leq 4N^2 \max_{1 \leq j \leq N} \max(|x_k|, |y_k|, |\theta_k - \tilde{\theta}_k|) \leq 4N^2 c,
\end{aligned}
$$

where the third and fourth lines are both applications of Lemma 3.1. Applying Lemma 3.1 once more, we have

$$
\begin{aligned}
\| \mathrm{dL}_{f(y)(f(x))^{-1}} \mathrm{d}f_x(h) - df_y(h) \|_{\mathrm{HS}} &\leq \|f(y) - f(x)\|_{\mathrm{HS}} + \| \mathrm{d}f_x(h) - \mathrm{d}f_y(h) \|_{\mathrm{HS}}, \\
&\leq 8N^2 c
\end{aligned}
$$

the second part of inequality (3.5). This completes the proof. □

Proof of Lemma 3.4. Let $x = (x_1, \ldots, x_N) \in [-c, c]^N$. We calculate:

$$
\begin{aligned}
&\| f(x) - f(0) \exp(\mathrm{dL}_{f(0)^{-1}} \mathrm{d}f_0(x)) \|_{\mathrm{HS}} \\
&= \| \prod_{k=1}^{N} \mathrm{R}_k e^{(\theta_k + x_k)a_k} \\
&\qquad - \prod_{k=1}^{N} \mathrm{R}_k e^{\theta_k a_k} \exp((\prod_{k=1}^{N} \mathrm{R}_k e^{\theta_k a_k})^{-1} \sum_{i=1}^{N} \prod_{k=1}^{N} \mathrm{R}_k \, e^{(\theta_k + x_k)a_k} (\mathrm{Id} + (h_k a_k - \mathrm{Id}) 1_{k=i})) \|_{\mathrm{HS}} \\
&= \| \prod_{k=1}^{N} \mathrm{R}_k e^{(\theta_k + x_k)a_k} \\
&\qquad - \prod_{k=1}^{N} \mathrm{R}_k e^{\theta_k a_k} \sum_{u=0}^{\infty} \frac{1}{u!} ((\prod_{k=1}^{N} \mathrm{R}_k e^{\theta_k a_k})^{-1} \sum_{i=1}^{N} \prod_{k=1}^{N} \mathrm{R}_k \, e^{(\theta_k + x_k)a_k} (\mathrm{Id} + (h_k a_k - \mathrm{Id}) 1_{k=i}))^u \|_{\mathrm{HS}} \\
&= \| \prod_{k=1}^{N} \mathrm{R}_k e^{\theta_k a_k} \sum_{u=0}^{\infty} \frac{(x_k a_k)^u}{u!} \\
&\qquad - \prod_{k=1}^{N} \mathrm{R}_k e^{\theta_k a_k} \sum_{u=0}^{\infty} \frac{1}{u!} ((\prod_{k=1}^{N} \mathrm{R}_k e^{\theta_k a_k})^{-1} \sum_{i=1}^{N} \prod_{k=1}^{N} \mathrm{R}_k \, e^{(\theta_k + x_k)a_k} (\mathrm{Id} + (h_k a_k - \mathrm{Id}) 1_{k=i}))^u \|_{\mathrm{HS}} \\
&\leq 8N^2 \max_{1 \leq k \leq N} |x_k|^2 \\
&\qquad + \| \sum_{i=1}^{N} \prod_{k=1}^{N} \mathrm{R}_k e^{\theta_k a_k} (\mathrm{Id} + (h_k a_k - \mathrm{Id}) 1_{k=i}) - \sum_{i=1}^{N} \prod_{k=1}^{N} \mathrm{R}_k \, e^{\theta_k a_k} (\mathrm{Id} + (h_k a_k - \mathrm{Id}) 1_{k=i}) \|_{\mathrm{HS}} \\
&= 8N^2 \max_{1 \leq k \leq N} |x_k|^2 \leq 8N^2 c^2,
\end{aligned}
$$

where the second-last line relies on the triangle inequality and repeated application of Lemma 3.1 to remove all terms that are of second or higher order in $\{a_k\}_{1 \leq k \leq N}$. This completes the proof of the lemma. □

Proof of Lemma 3.5. By Equation (3.4),

$$\langle \mathrm{d}f_x(h), \mathrm{d}f_x(h')\rangle_{\mathrm{HS}} = \mathrm{Tr}[\mathrm{d}f_x(h)\mathrm{d}f_x(h')^\dagger] \tag{A.1}$$
$$= \mathrm{Tr}\Big[\Big(\sum_{i=1}^N \prod_{k=1}^N \mathrm{R}_k\, e^{(\theta_k+x_k)a_k}(\mathrm{Id} + (h_k a_k - \mathrm{Id})1_{k=i})\Big)\Big(\sum_{i=1}^N \prod_{k=1}^N \mathrm{R}_k\, e^{(\theta_k+x_k)a_k}(\mathrm{Id} + (h'_k a_k - \mathrm{Id})1_{k=i})\Big)^\dagger\Big]$$
$$= \sum_{i=1}^N \mathrm{Tr}\Big[\prod_{k=1}^N \mathrm{R}_k\, e^{(\theta_k+x_k)a_k}(\mathrm{Id} + (h_k a_k - \mathrm{Id})1_{k=i}) \prod_{k=N}^1 (\mathrm{Id} - (h'_k a_k - \mathrm{Id})1_{k=i}) e^{-(\theta_k+x_k)a_k}\mathrm{R}_k^{-1}\Big]$$
$$+ \sum_{1\le i\ne j\le N} \mathrm{Tr}\Big[\prod_{k=1}^N \mathrm{R}_k\, e^{(\theta_k+x_k)a_k}(\mathrm{Id} + (h_k a_k - \mathrm{Id})1_{k=i}) \prod_{k=N}^1 (\mathrm{Id} - (h'_k a_k - \mathrm{Id})1_{k=j}) e^{-(\theta_k+x_k)a_k}\mathrm{R}_k^{-1}\Big]$$
$$\equiv \sum_{i=1}^N S_i + \sum_{1\le i\ne j\le N} S_{ij}.$$

We calculate terms of the form S_j and S_{ij} separately. For any S_i, applying the cyclic permutation property of the trace operator yields

$$S_i = \mathrm{Tr}\Big[\prod_{k=1}^N \mathrm{R}_k\, e^{(\theta_k+x_k)a_k}(\mathrm{Id} + (h_k a_k - \mathrm{Id})1_{k=i}) \prod_{k=N}^1 (\mathrm{Id} - (h'_k a_k - \mathrm{Id})1_{k=i}) e^{-(\theta_k+x_k)a_k}\mathrm{R}_k^{-1}\Big]$$
$$= \mathrm{Tr}\Big[\mathrm{R}_i(h_i a_i)(-h'_i a_i)\mathrm{R}_i^{-1}] = -h_i h'_i \mathrm{Tr}[\mathrm{R}_i a_i^2 \mathrm{R}_i^{-1}\Big] = -h_i h'_i \mathrm{Tr}[a_i^2] = h_i h'_i. \tag{A.2}$$

For any S_{ij} with $i < j$,

$$S_{ij} = \mathrm{Tr}\Big[\prod_{k=1}^N \mathrm{R}_k\, e^{(\theta_k+x_k)a_k}(\mathrm{Id} + (h_k a_k - \mathrm{Id})1_{k=i}) \prod_{k=N}^1 (\mathrm{Id} - (h'_k a_k - \mathrm{Id})1_{k=j}) e^{-(\theta_k+x_k)a_k}\mathrm{R}_k^{-1}\Big]$$
$$= -h_i h'_j \mathrm{Tr}[a_i M_{ij} \mathrm{R}_j\, e^{(\theta_j+x_j)a_j} a_j\, e^{-(\theta_j+x_j)a_j}\mathrm{R}_j^{-1} M_{i,j}^{-1}]$$
$$= -h_i h'_j \mathrm{Tr}[a_i M_{i,j} \mathrm{R}_j a_j \mathrm{R}_j^{-1} M_{i,j}^{-1}]$$
$$= h_i h'_j D_{ij}. \tag{A.3}$$

For $j < i$, a similar calculation gives

$$S_{ij} = h_i h'_j D_{i,j}.$$

Combining Equalities (A.1), (A.2) and (A.3) completes the proof. □

Proof of Lemma 3.6. We note that f is clearly smooth, and so we must only check that it is bijective. This result will follow almost immediately from Lemma 10 of [Ray02] and our bounds in Lemmas 3.4 and 3.5.

We begin to set up notation. For a point x in a metric space (Ω, d) and constant $\delta > 0$, let $\mathcal{B}_\delta(x) = \{y \in \Omega : d(x,y) \le \delta\}$ be the ball of radius δ around x. We then define a map $\mathcal{F}$ from $\mathcal{B}_{n^{-6}\phi_n}(\mathrm{Id}) \subset \mathrm{SO}(n)$ to $\mathbb{R}^N$ as follows. Let $x \in \mathcal{B}_{n^{-6}\phi_n}(\mathrm{Id})$. Since the exp map is surjective and sends lines to geodesic curves, we can write $x = \exp(h)$ for some $h \in \mathfrak{so}(n)$ with $\|h\|_{\mathrm{HS}} \le 2n^{-6}\phi_n$. Furthermore, we can write $h = \sqrt{2}\sum_{i=1}^N h_i a_i$. We then define $\mathcal{F}(x) = (h_1, h_2, \ldots, h_N)$. Finally, we define the map $g = \mathcal{F} \circ f : [-c,c]^N \mapsto \mathbb{R}^N$.

We point out that $\mathcal{F}$ has small distortion: for $x, y \in \mathrm{SO}(n)$ with $\mathcal{F}(x) = h_x$, $\mathcal{F}(y) = h_y$,

$$\|x - y\|_{\mathrm{HS}} = \|\exp(h_x) - \exp(h_y)\|_{\mathrm{HS}} = \|h_x - h_y\| + O(N^2\|h_x - h_y\|^2). \tag{A.4}$$

We now obtain the estimates required to use Lemma 10 of [Ray02]. Following the notation of that paper, we set $\rho = \frac{1}{256}n^{-6}\phi_n$, $\delta = \frac{\phi_n}{8}$ and $\rho_* = \frac{\delta}{8}\rho$. For $x, z \in [-c, c]^N$ with $\|x - z\| = \rho$,

$$\begin{aligned}
\|g(x) - g(z)\|_{\mathrm{HS}} &= \|g(x) - \mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(x))) + \mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(x))) \\
&\quad + \mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(z))) - \mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(z))) - g(z)\|_{\mathrm{HS}} \\
&\geq \|\mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(x))) - \mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(z)))\|_{\mathrm{HS}} \\
&\quad - \|\mathcal{F}(f(x)) - \mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(x)))\|_{\mathrm{HS}} \\
&\quad - \|\mathcal{F}(f(z)) - \mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(z)))\|_{\mathrm{HS}} \\
&\geq \|\mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(x))) - \mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(z)))\|_{\mathrm{HS}} \\
&\quad - 2\|f(x) - f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(x))\|_{\mathrm{HS}} \\
&\quad - 2\|f(z) - f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(z))\|_{\mathrm{HS}} \\
&\geq \|f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(x)) - f(0)\exp(\mathrm{d}L_{f(0)^{-1}}\mathrm{d}f_0(z))\|_{\mathrm{HS}} - 32N^2c^2,
\end{aligned} \tag{A.5}$$

where the second-last inequality follows from inequality (A.4) and Lemma 3.4, and the last inequality is due to Lemma 3.4. By Inequalities (3.8) and (A.4),

$$\begin{aligned}
&\|\mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(x))) - \mathcal{F}(f(0)\exp(\mathrm{dL}_{f(0)^{-1}}\mathrm{d}f_0(z)))\|_{\mathrm{HS}} \\
&\qquad \geq \frac{\phi_n}{4}\|x - z\| - O(N^3\|x - z\|^2).
\end{aligned} \tag{A.6}$$

Combining Inequalities (A.5) and (A.6), we conclude

$$\|g(x) - g(z)\|_{\mathrm{HS}} \geq \frac{\phi_n}{8}\|x - z\|.$$

This proves the first condition of Lemma 10 of [Ray02]: $\{\|x - z\|_{\mathrm{HS}} > \rho\}$ implies that $\{\|g(x) - g(z)\| > \delta\rho\}$. The second condition of Lemma 10 of [Ray02] follows immediately from the second part of inequality (3.8). Thus, g satisfies the requirements of Lemma 10 of [Ray02] with ρ, ρ_*, δ as above, and so g is an injective map. But this implies that f is injective as well, completing the proof. □

A.2. **Probability Estimates.**

Proof of Lemma 3.7. We recognize that s_N is exactly the time required to collect all N coupons in the standard 'coupon collector problem' with N coupons. Our Equation (3.11) above is Equation (2) of [ER61]. □

Proof of Lemma 3.8. We recognize that $s_N - \mathcal{Q}N^2\lceil\log(N)\rceil$ has exactly negative binomial distribution with both parameters equal to N. The bound is then a standard tail bound for the negative binomial distribution (see, *e.g.*, the calculation in [Bro]). □

Proof of Lemma 3.9. For fixed $0 \leq t \leq n^5$, let $\mathcal{E}(t)$ be the event $\{\sup_{0\leq s\leq t}\|\hat{X}_s - \hat{Y}_s\|_{\mathrm{HS}} \leq \|\hat{X}_0 - \hat{Y}_0\|_{\mathrm{HS}}n^{5+C}\}$. By the main calculation on pp. 1216 of [Oli09], we have for all $0 \leq t \leq n^5$ that

$$\mathbb{E}[\|\hat{X}_{t+1} - \hat{Y}_{t+1}\|_{\mathrm{HS}}\mathbf{1}_{\mathcal{E}(t)}] \leq \|X_0 - Y_0\|_{\mathrm{HS}}(1 + \sqrt{\|\hat{X}_0 - \hat{Y}_0\|_{\mathrm{HS}}}n^{7+C})^t.$$

Inequality (3.12) follows immediately from Markov's inequality and a union bound over $1 \le t \le n^5$.

Inequality (3.13) follows from noting that,

$$\begin{aligned}
\|X_t - \hat{X}_t\|_{\mathrm{HS}} &= \|\prod_{s=0}^{t-1} \mathrm{R}(i_s(x), \eta_s(x)) - \prod_{s=0}^{t-1} \mathrm{R}(i_s(x), \theta_s(x))\|_{\mathrm{HS}} \\
&\le \sum_{u=0}^{t-1} \|(\mathrm{R}(i_u(x), \eta_t(x)) - \mathrm{R}(i_u(x), \theta_u(x))) \prod_{s=u+1}^{t-1} \mathrm{R}(i_s(x), \eta_s(x))\|_{\mathrm{HS}} \\
&= \sum_{u=0}^{t-1} \|(\mathrm{R}(i_u(x), \eta_u(x)) - \mathrm{R}(i_u(x), \theta_u(x)))\|_{\mathrm{HS}} \\
&\le \sum_{u=0}^{t-1} 6\|\eta_u(x) - \theta_u(x)\| \\
&\le 6t\epsilon_n \le 6n^5\epsilon_n,
\end{aligned}$$

where we use Lemma 3.1 in the second line. Inequality (3.14) can be proved analogously. □